

Part of 45
Pages of 9



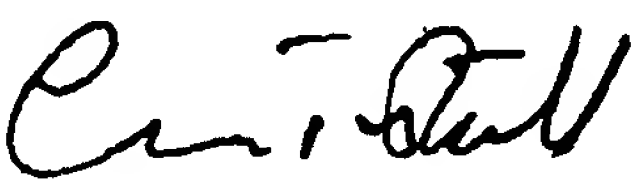
Notice of Opposition to a European Patent

To the
European Patent Office

Tabulation marks		for EPO use only	
I. Patent opposed Patent No. EP 1022642 B1 Application No. 00108822.8 Date of mention of the grant in the European Patent Bulletin (Art. 97(4), 98(1) EPC) 05-09-2001		Opp. No. OPPO (1)	
Title of the invention: INTEGRATED CIRCUIT I/O USING A HIGH PERFORMANCE BUS INTERFACE			
II. Proprietor of the Patent RAMBUS INC. first named in the patent specification			
III. Opponent Opponent's or representative's reference (max. 15 characters)		OREF	
Name Address State of residence or of principal place of business Telephone/Telex/Fax Multiple opponents		OPPO (2) MICRON EUROPE Ltd Micron House Wellington Business Park Dukes Ride Crowthorne Berkshire RG45 6LS UNITED KINGDOM UNITED KINGDOM <input type="checkbox"/> further opponents see additional sheet	
IV. Authorisation 1. Representative (Name only one representative to whom notification is to be made) Name Address of place of business Telephone/Telex/Fax Additional representative(s) 2. Employee(s) of the opponent authorised for these opposition proceedings under Art. 133(3) EPC Authorisation(s) To 1./2.		OPPO (3) Christopher Stephen Tunstall Harrison Goddard Foote Tower House Merrion Way Leeds LS2 8PA UNITED KINGDOM +44 113 290 1400 +44 113 244 2829 <input type="checkbox"/> (on additional sheet/see authorisation) OPPO (5) Name(s): <input checked="" type="checkbox"/> not considered necessary <input type="checkbox"/> has/have been registered under No. <input type="checkbox"/> is/are enclosed	

		for EPO use only
V. Opposition is filed against — the patent as a whole <input checked="" type="checkbox"/> — claim(s) No(s). <input type="checkbox"/> Claims 1-17		
VI. Grounds for opposition: Opposition is based on the following grounds: (a) the subject-matter of the European patent opposed is not patentable (Art. 100(a) EPC) because: — it is not new (Art. 52(1); 54 EPC) <input checked="" type="checkbox"/> — it does not involve an inventive step (Art. 52(1); 56 EPC) <input checked="" type="checkbox"/> — patentability is excluded on other grounds, i. e. <input type="checkbox"/> Art. <input type="text"/> (b) the patent opposed does not disclose the invention in a manner sufficiently clear and complete for it to be carried out by a person skilled in the art (Art. 100(b) EPC; see Art. 83 EPC). <input checked="" type="checkbox"/> (c) the subject-matter of the patent opposed extends beyond the content of the application/ of the earlier application as filed (Art. 100(c) EPC; see Art. 123(2) EPC). <input checked="" type="checkbox"/>		
VII. Facts and arguments (Rule 55(c) EPC) presented in support of the opposition are submitted herewith on a separate sheet (annex 1)		<input checked="" type="checkbox"/>
VIII. Other requests: That the patent be revoked in its entirety.		

IX. Evidence presented		for EPO use only
<p>Undated <input type="checkbox"/></p> <p>will be filed at a later date <input checked="" type="checkbox"/></p>		
A. Publications:		Publication date
1 See Statement of Grounds of Opposition		
Particular relevance (page, column, line, fig.):		
2		
Particular relevance (page, column, line, fig.):		
3		
Particular relevance (page, column, line, fig.):		
4		
Particular relevance (page, column, line, fig.):		
5		
Particular relevance (page, column, line, fig.):		
6		
Particular relevance (page, column, line, fig.):		
7		
Particular relevance (page, column, line, fig.):		
Continued on additional sheet <input type="checkbox"/>		
B. Other evidence		
Continued on additional sheet <input type="checkbox"/>		

		for EPO use only																																	
X. Payment of the opposition fee is made <input checked="" type="checkbox"/> as indicated in the enclosed voucher for payment of fees and costs (EPO Form 1010) <input type="checkbox"/>																																			
XI. List of documents <table border="1"><thead><tr><th>Enclosure No.</th><th></th><th>No. of copies</th></tr></thead><tbody><tr><td>0</td><td><input checked="" type="checkbox"/> Form for notice of opposition</td><td><input type="text" value="2"/> (min. 2)</td></tr><tr><td>1</td><td><input checked="" type="checkbox"/> Facts and arguments (see VII.)</td><td><input type="text" value="2"/> (min. 2)</td></tr><tr><td>2</td><td colspan="2">Copies of documents presented as evidence (see IX.)</td></tr><tr><td>2a</td><td><input type="checkbox"/> -- Publications</td><td><input type="text"/> (min. 2 of each)</td></tr><tr><td>2b</td><td><input type="checkbox"/> -- Other documents</td><td><input type="text"/> (min. 2 of each)</td></tr><tr><td>3</td><td><input type="checkbox"/> Signed authorisation(s) (see IV.)</td><td><input type="text"/></td></tr><tr><td>4</td><td><input checked="" type="checkbox"/> Voucher for payment of fees and costs (see X.)</td><td><input type="text" value="1"/></td></tr><tr><td>5</td><td><input type="checkbox"/> Cheque</td><td><input type="text"/></td></tr><tr><td>6</td><td><input type="checkbox"/> Additional sheet(s)</td><td><input type="text"/> (min. 2 of each)</td></tr><tr><td>7</td><td><input type="checkbox"/> Other (please specify here):</td><td><input type="text"/></td></tr></tbody></table>		Enclosure No.		No. of copies	0	<input checked="" type="checkbox"/> Form for notice of opposition	<input type="text" value="2"/> (min. 2)	1	<input checked="" type="checkbox"/> Facts and arguments (see VII.)	<input type="text" value="2"/> (min. 2)	2	Copies of documents presented as evidence (see IX.)		2a	<input type="checkbox"/> -- Publications	<input type="text"/> (min. 2 of each)	2b	<input type="checkbox"/> -- Other documents	<input type="text"/> (min. 2 of each)	3	<input type="checkbox"/> Signed authorisation(s) (see IV.)	<input type="text"/>	4	<input checked="" type="checkbox"/> Voucher for payment of fees and costs (see X.)	<input type="text" value="1"/>	5	<input type="checkbox"/> Cheque	<input type="text"/>	6	<input type="checkbox"/> Additional sheet(s)	<input type="text"/> (min. 2 of each)	7	<input type="checkbox"/> Other (please specify here):	<input type="text"/>	
Enclosure No.		No. of copies																																	
0	<input checked="" type="checkbox"/> Form for notice of opposition	<input type="text" value="2"/> (min. 2)																																	
1	<input checked="" type="checkbox"/> Facts and arguments (see VII.)	<input type="text" value="2"/> (min. 2)																																	
2	Copies of documents presented as evidence (see IX.)																																		
2a	<input type="checkbox"/> -- Publications	<input type="text"/> (min. 2 of each)																																	
2b	<input type="checkbox"/> -- Other documents	<input type="text"/> (min. 2 of each)																																	
3	<input type="checkbox"/> Signed authorisation(s) (see IV.)	<input type="text"/>																																	
4	<input checked="" type="checkbox"/> Voucher for payment of fees and costs (see X.)	<input type="text" value="1"/>																																	
5	<input type="checkbox"/> Cheque	<input type="text"/>																																	
6	<input type="checkbox"/> Additional sheet(s)	<input type="text"/> (min. 2 of each)																																	
7	<input type="checkbox"/> Other (please specify here):	<input type="text"/>																																	
XII. Signature of opponent or representative Place Leeds, UNITED KINGDOM Date September 10, 2001  Christopher Stephen Tunstall <small>Please type name under signature. In the case of legal persons, the position which the person signing holds within the company should also be typed.</small>																																			

EUROPEAN PATENT NO. 1,022,642

STATEMENT OF GROUNDS OF OPPOSITION

1. **PRIMA FACIE DIVISIONAL STATUS**

- 1.1. The Patent has been granted pursuant to an application, no. 00108822.8 ("The Application"), divided from an earlier application, no. 99118308.8 ("The Parent Application"), itself a division of a still earlier application, no. PCT/US91/02590 ("The Grandparent Application"). This is not permitted. As discussed in decision T904/97, the European Patent Convention does not explicitly foresee the possibility of an application being divided from an earlier divisional application. The *Travaux Préparatoires*, on the other hand, contain a reference to sequences of divisional applications that evidences an intention not to permit the filing of sequences of divisional applications. According to the minutes of its 12th meeting held from 26 February to 6 March 1964 in Brussels, the EEC Patents Working Party rejected a proposal relating to "a system of divisions in a cascade, which could have constituted a dilatory manoeuvre." In the working party's opinion, this was ruled out by the division system provided for in Article 80 of the Preliminary Draft Convention (see the document 2632/IV/64-F, 15 April 1964, page 31). The Application was never a divisional application at all. It was a regular application having a filing date of April 26, 2000 and no valid priority claim.

2. **ADDED SUBJECT-MATTER ART. 76(1) EPC**

- 2.1. The following discussion assumes that the Patent has been granted pursuant to a procedurally valid divisional filing. According to Art. 76(1), it was not permissible to include subject-matter in the Parent Application that was not included in the Grandparent Application. Likewise, it was not permissible to include subject-matter in the Application that was not included in the Parent Application. Therefore, it was not permissible to include subject-matter in the Application that was not included in the Grandparent Application. For present purposes, it is unnecessary to determine whether matter was added when the Parent Application was filed or when the Application was filed or later and it is instructive to look at the overall picture by comparing the Patent with the Grandparent Application.
- 2.2. The Patent discloses subject-matter not disclosed in the Grandparent Application as originally filed on two counts, namely impermissible claim broadening and impermissible intermediate claim generalisations.
- 2.3. **Impermissible Claim Broadening**
- 2.3.1. The Grandparent Application contained 7 objectives of the invention and 21 independent claims. The interrelationships between the objectives and claims need not be explored here in full. For present purposes, it is clear that granted claim 1 is based upon PCT claim 103, this being the broadest claim to modifiable access times.
- 2.3.2. One objective of the invention as set out in the Grandparent Application survived in the Patent, albeit in modified form. It was this.
- 2.3.2.1. "One object of the present invention is to use a new bus interface built into semiconductor devices to support high-speed access to large blocks of data from a single memory device by an external user of the data, such as a microprocessor, in an efficient and cost-effective manner" [Grandparent Application, 6:8-12]. This will be referred to as "objective 1."
- 2.3.3. The Grandparent Application included a further objective which, as discussed below, is of relevance.

2.3.3.1. "Another object of this invention is to provide devices, especially DRAMs, suitable for use with the bus architecture of the invention." [Grandparent Application 7:5-7]. This will be referred to as "objective 2."

2.3.4. Objective 1 was not achieved by PCT claim 103.

2.3.4.1. There are few references in the Grandparent Application to the advantages of modifiable access times. One reference however reads as follows.

2.3.4.2. "A request packet and the corresponding bus access are separated by a selected number of bus cycles, allowing the bus to be used in the intervening bus cycles by the same or other masters for additional requests or brief bus accesses. Thus multiple, independent accesses are permitted, allowing maximum utilisation of the bus for transfer of short blocks of data. Transfers of long blocks of data use the bus efficiently even without overlap because the overhead due to bus address, control and access times is small compared to the total time to request and transfer the block" [Grandparent Application, 15:23-16:7].

2.3.4.3. This passage is a very clear statement of the advantages of modifiable access times. In the transfer of short blocks of data, modifiable access times allow for interleaving of requests and thus greater efficiency. In the transfer of long blocks of data, the need to transfer access time information on the bus is an overhead. It reduces efficiency. There is no possibility of interleaving long block requests. The preferred embodiments described in the Grandparent Application do not allow it. However, this overhead is an acceptable efficiency reduction because, compared with the overall length of the transaction, the overhead is small.

2.3.4.4. Objective 1 concerns "high-speed access to large blocks of data ... in an efficient and cost-effective manner." Modifiable access times contribute nothing to the achievement of that objective; as acknowledged in the Grandparent Application itself, they detract from it. PCT claim 103 does not achieve this objective.

2.3.4.5. Objective 1 was really achieved by the measures described in "DRAM Column Access Modification" [Grandparent Application, 59:4-62:2]. These measures, including internal I/O multiplexing, allow the interfacing of a memory device running at a relatively slow internal clock rate with the high-speed bus of the invention [Grandparent Application, 60:1-6]. This increases the bandwidth of DRAM access [Grandparent Application, 59:19-2; 61:3-13]. It was acknowledged that the invention lay not in these measures *per se*, but in their use with the high-speed highly multiplexed bus of the invention [Grandparent Application 59:23-25]. These measures were the subject-matter of claims 82-90 and 114-123 of the Grandparent Application.

2.3.5. Objective 2 was achieved by PCT claim 103.

2.3.5.1. There can be no doubt whatsoever that in objective 2 "the bus architecture of the invention" meant a bus architecture in which the bus includes a plurality of bus lines for carrying substantially all address, data and control information needed by the device for communication with substantially every other device connected to the bus, and has substantially fewer bus lines than the number of bits in a single address [Grandparent Application, 7:10-19; 11:16-12:10]. A bus including a plurality of such general purpose lines, carrying in a time-multiplexed manner substantially all address, data and control information needed by the device for communication with substantially every other device connected to the bus, and having substantially fewer bus lines than the number of bits in a single address will be referred to in the present document as a "highly multiplexed bus." This definition is constantly reinforced throughout the application including the summary of the invention and the beginning of the specific description [Grandparent Application, 7:10-19; 7:25-8:2; 11:16-25].

2.3.5.2. In all, claims 82-90 and 95-150 of the Grandparent Application achieve this objective: it is a wide objective, but it is restricted by reference to the bus architecture of the invention. Each of these claims requires a device that can be used in a bus architecture where the bus includes a plurality of general purpose lines, carrying in a time-multiplexed manner substantially all address, data and control information needed by the device for communication with substantially every other device connected to the bus, and having substantially fewer bus lines than the number of bits in a single address, that is to say, a highly multiplexed bus.

2.3.6. In summary, in relation to the granted claims, the invention as originally disclosed was that disclosed in PCT claim 103, in the context of the overall description and objective 2.

2.3.7. The Technical Board of Appeal case law on claim broadening is settled.

2.3.7.1. The test for whether claim broadening is contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also, was clearly enunciated in decision T331/87, following decisions T194/84 and T260/85. The test was approved in decision T514/88 and is threefold [T331/87, Reasons:6].

2.3.7.2. "The removal of an integer from a claim may not violate Art. 123(2) EPC provided the skilled reader would directly and unambiguously recognise that:

1. the integer was not explained as essential in the original disclosure;
2. it is not, as such, indispensable for the function of the invention in the light of the technical problem it serves to solve; and
3. the replacement or removal requires no real modification of other integers to compensate for the change."

2.3.8. Granted claim 1 lacks an essential element of the invention as originally disclosed.

2.3.8.1. Granted claim 1 requires a semiconductor memory device including output drivers for "outputting data." PCT claim 103 explicitly required the device to interface to a highly multiplexed bus. If the removal of this explicit requirement means that claim 1 does not require the data to be output onto a highly multiplexed bus, this is a crucial distinction because the requirement of PCT claim 103 for a highly multiplexed bus implied that the device claimed must have certain interface circuits that would allow it to demultiplex and decode relevant bus transactions. Therefore, granted claim 1, if it no longer contains this implied requirement, is broader than PCT claim 103. Such broadening is inadmissible.

2.3.8.2. Was compatibility with the highly multiplexed bus explained as essential in the original disclosure?

2.3.8.2.1. According to decisions T260/85 and T527/88, in determining what is explained or disclosed in the application, it is necessary to read the document as a whole and individual passages within it in context. In view of this, a line-by-line analysis is no substitute for gaining an overall impression of the document by reading it from start to finish. However, the following points are mentioned to give a flavour for the disclosure of the Grandparent Application.

2.3.8.2.1.1. The application was entitled "Integrated Circuit I/O Using a High Performance Bus Interface."

- 2.3.8.2.1.2. The discussion of the prior art drew many distinctions between the numerous documents mentioned and the invention, some of which were applicable to different inventions from the inventions originally claimed. However, in every case but one, the documents were distinguished from the invention on the basis that it did not possess features of the bus interface [Grandparent Application, 3:14-18; 3:25-4:3; 4:5-7; 4:13-14; 4:16-18; 5:13-14; 6:5-7]. The single exception was a reference that described a clocking scheme that was distinguished from the clocking scheme used in the invention [Grandparent Application, 5:19-25]. The only claims in the Grandparent Application that did not contain a detailed description of the highly multiplexed bus were either directed to this clocking scheme, claims 73-81, or the packaging of the device, claims 91-94.
- 2.3.8.2.1.3. Moreover, the state of the art was summed up thus: "None of the buses described in patents or other literature use only bused connections. All contain some point-to-point connections on the backplane" [Grandparent Application 5:13-15].
- 2.3.8.2.1.4. The summary of the invention opens with a discussion of the highly multiplexed bus [Grandparent Application, 7:10-19].
- 2.3.8.2.1.5. The objective achieved by PCT claim 103 was concerned with the highly multiplexed bus. This is highly relevant [T514/88].
- 2.3.8.2.1.6. The entire specific description is concerned with a highly multiplexed bus.
- 2.3.8.2.1.7. All of the original claims, apart from those directed to the clocking scheme or the packaging, contained a detailed description of the highly multiplexed bus.
- 2.3.8.2.1.8. Furthermore, even the discussion of the advantages of modifiable access times, as set out in para. 1.1.2.2 above, can be seen to presuppose the highly multiplexed bus. As discussed above, the inventors considered modifiable access times to be advantageous when using the bus for the transfer of small blocks of data, because in the bus cycles intervening between a request packet and the satisfaction of the request, the bus could be used for additional requests or brief bus accesses [Grandparent Application, 15:23-16:7]. In the highly multiplexed system, where address, data and control signals are carried on the same lines, it is of course important to ensure that data resulting from the satisfaction of an earlier request does not collide on the bus with address and control information relating to a subsequent request. In a non-multiplexed system where address, data and control signals are carried on separate lines, data resulting from the satisfaction of an earlier request cannot collide with address and control information relating to a subsequent request. Modifiable access times are disclosed only as a consequence of the highly multiplexed bus and do not in the inventors' view have any existence independent of it.
- 2.3.8.2.2. All of these matters point to the essentiality of the highly multiplexed bus. Indeed, the whole tenor of the Grandparent Application was that it concerned a new bus architecture. There is no disclosure, suggestion or implication that anything else was contemplated. It is what the inventors contemplated, as objectively determined from their original application, that counts [T260/85, Reasons:10]. The Grandparent Application clearly explained that compatibility with the highly

multiplexed bus was an essential element of the invention of PCT claim 103. No other reading or explanation is possible.

2.3.8.3. Was compatibility with the highly multiplexed bus, as such, indispensable for the function of the invention in the light of the technical problem it serves to solve?

2.3.8.3.1. As is clear from decision T331/87 and T514/88 which followed it, the "technical problem" to be considered here must at the very least include the relevant objective of the invention as originally disclosed. It cannot be some redefined technical problem such as might be used to analyse inventive step, because such technical problems are defined vis-à-vis the objectively most relevant of the known prior art and therefore take into account matter not disclosed in the application as filed. Even the achievement of advantages that are disclosed in the original application, but not set out as objectives of the invention, will not do. They are relevant only to intermediate generalisation, as is discussed below in paragraph 1.4.1.3.3. The purpose of the test set out in T331/87 is to determine objectively which integers of an original claim are essential to the invention. Clearly, if any such integer is not required to achieve the stated objective of the invention, that is a relevant fact. On the other hand, if any such integer is not required to achieve an entirely different objective, with which the invention in its broadest sense was not concerned, that fact is equally clearly of no consequence whatsoever.

2.3.8.3.2. As discussed in paragraphs 2.4.2 – 2.4.5 above, notwithstanding the appearance in the Patent of one objective, PCT claim 103 was concerned with another of the original objectives, objective 2. Thus, the technical problem that PCT claim 103 served to solve was to provide devices, especially DRAMs, suitable for use with the highly multiplexed bus architecture in question. Therefore, compatibility with the highly multiplexed bus was not just indispensable; it was a prerequisite.

2.3.8.3.3. Moreover, as explained above, not only was the objective to be achieved expressly stated to concern compatibility with the highly multiplexed bus, but also the advantages described as attributable to modifiable access times, are discussed wholly within the context of the packet-based bus [Grandparent Application, 15:23-24]. The packet-based bus here described is, of course, the highly multiplexed bus.

2.3.8.4. Would the removal of the requirement for compatibility with the highly multiplexed bus require any real modification of other integers to compensate for the change?

2.3.8.4.1. Again, this is a straightforward question to answer. Removal of this requirement has far-reaching consequences.

2.3.8.4.2. Firstly, the programmable access time register would no longer need to be accessible via the bus. This was an explicit requirement of PCT claim 103 and stemmed from the control information being multiplexed onto the bus. Thus, the circuitry of the device that allows the access time register to be programmed would be able to take its inputs from non-bused signals and need not demultiplex them from bused signals, as is the case with PCT claim 103.

2.3.8.4.3. Secondly, the nature of the "request" of PCT claim 103 is defined by the highly multiplexed nature of the bus. Removal of this interdependence from granted claim 1 would be a radical departure from the teaching of the application as filed, in which requests had to be

encoded and multiplexed onto the bus lines in the way similar to that shown in and described with reference to Fig. 4 of the Grandparent Application [Grandparent Application, 21:21-24:2]. Any removal of the requirement for compatibility with the highly multiplexed bus would radically alter the meaning of the term "request" as between PCT claim 103 and granted claim 1.

2.3.8.4.4. Thirdly, as discussed, modifiable access times are disclosed as a consequence of the highly multiplexed bus and do not in the inventors' view have an existence independent of it. Any removal of the requirement for compatibility with the highly multiplexed bus would also remove the justification for modifiable access times as perceived by the inventors.

2.3.8.5. For all these reasons, if the requirement in PCT claim 103 for compatibility with the highly multiplexed bus is absent from granted claim 1, its absence fails all three limbs of the test established by the Technical Board of Appeal. It would amount to inadmissible claim broadening contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.4. It follows for precisely the reasons discussed above that neither granted claim 1 nor any of the remaining claims is entitled to the priority date claimed.

2.5. Impermissible Intermediate Generalisation

2.5.1. The Technical Board of Appeal case law on intermediate generalisation is settled.

2.5.1.1. The test for whether intermediate generalisation is contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also, was clearly enunciated in decision T284/94, following decision T17/86.

2.5.1.2. "An amendment of a claim by the introduction of a technical feature taken in isolation from the description of a specific embodiment is not allowable under Art. 123(2) EPC if it is not clear beyond any doubt for a skilled reader from the application documents as filed that the subject-matter of the claim thus amended provides a complete solution to a technical problem unambiguously recognisable from the application."

2.5.1.3. A number of other passages from this decision shed useful light on how this test is to be applied, in particular the following statements.

2.5.1.3.1. Referring to T17/86, an isolated technical feature "may be introduced into a claim without contravening Art. 123(2) EPC, provided that it is "evident beyond doubt to a skilled person reading the description that this isolated technical feature on its own enables the object in view to be achieved" " [T284/94, Reasons point 2.1.3, para. 2].

2.5.1.3.2. "In following this decision, the object to be achieved by the subject-matter of the amended claim has to be established as well as whether the claims define all means necessary for achieving this object" [T284/94, Reasons point 2.1.4, para. 1].

2.5.1.3.3. "Because of the fact that features disclosed in the context of a specific embodiment and added to a claim may achieve in an unambiguously recognisable manner an object different from that present in the introductory part of a description, it should further be ascertained whether such a further object is disclosed and whether it is clear beyond doubt for a skilled person reading the application as filed that the added technical features on their own achieve this further object" [T284/94, Reasons point 2.1.4, para. 3]. This makes clear the important distinction

between objects that are relevant to claim broadening, i.e. the original objectives of the invention set out as such, and objects relevant to intermediate generalisation, which can be derived from other parts of the application, but only to the extent they are disclosed.

2.5.2. Granted claim 1 includes a technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.5.2.1. PCT claim 103 states that access time information "may be transmitted to said [access time] register via said bus which establishes a predetermined amount of time that said semiconductor device thereafter must wait before using said bus in response to a request." It is clear from the language of PCT claim 103, in particular the phrase "thereafter must wait ... in response to a request," that it applies to all types of request. This is consistent with the preferred embodiment, in which transaction requests contain access time information in all cases [Grandparent Application, 27:1-15]. This requirement can easily be understood. The advantage of a modifiable access time register is, as discussed above, to allow "maximum utilisation of the bus for transfer of short blocks of data" [Grandparent Application, 15:1-3]. Maximum utilisation of the bus can only be achieved if all the devices connected to the bus use a programmable access time register to time their response to all requests directed to them.

2.5.2.2. Granted claim 1 requires the semiconductor memory device to respond in a certain way to read requests only. If granted claim 1 were limited to read only devices such as ROMs, this limitation to read requests only would not be objectionable. However, granted claim 1 also covers read/write devices such as RAMs that have programmable access times for read requests only, and not for write requests. This represents an impermissible intermediate generalisation.

2.5.2.3. As discussed, the technical problem solved by modifiable access times, as claimed in PCT claim 103, was to maximise bus utilisation for exchange of small blocks of data in a highly multiplexed system [Grandparent Application, 16:1-3]. Modifiable access time DRAM read operations were discussed in the Grandparent Application in conjunction with modifiable access time DRAM write operations. Whereas bus utilisation may be improved by providing for modifiable access times in respect of read operations only, it is not a complete solution to the problem of maximising bus utilisation. For the achievement of that objective, the device must have programmable access times in respect of all operations that require it to utilise the bus in response to a request. Granted claim 1 is not so limited.

2.5.2.4. For this reason, the introduction into granted claim 1 of the requirement for modifiable access times in respect of read operations only fails the test established by the Technical Board of Appeal. It amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.5.3. Granted claim 1 includes a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.5.3.1. PCT claim 103 is not limited to semiconductor memory devices; it also covers devices such as disk controllers which, by definition, do not connect only to the bus [Grandparent Application 11:19-21]. A disk controller does not "send and receive substantially all address, data and control information over said bus", because it exchanges data with the disk drive. It does, however, exchange via the bus "substantially all address, data and control information needed for communication with substantially every other device connected to the bus."

- 2.5.3.2. All the semiconductor memory devices disclosed in the Grandparent Application communicate exclusively over the bus, including those disclosed in PCT claim 104. Semiconductor memory devices are not required to communicate with devices that are not connected to the bus. The broadest disclosure of a semiconductor memory device in the Grandparent Application is to be found in PCT claims 104, 107, 110, 112, 115, 117, 119, 122, 134 and 135. In each case without exception, the device is stated to have an exclusively bus-based interface.
- 2.5.3.3. The common objective to be achieved by the memory devices according to each of these claims, and indeed every single memory device disclosed in the Grandparent Application, was the same as that to be solved by PCT claim 103: to provide devices, especially DRAMs, suitable for use with the bus architecture of the invention. Semiconductor memory devices are just a special case. According to the Grandparent Application, memory devices for use in the bus architecture of the invention contain "only a single bus interface with no other signal pins" [Grandparent Application, 8:9-13]. It is clear that this common objective is only completely achieved if the semiconductor memory device has an exclusively bus-based interface.
- 2.5.3.4. Therefore, in accordance with the established case law of the Technical Board of Appeal, restricting any claim to a semiconductor memory device, but not to an exclusively bus-based interface, amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.
- 2.5.4. Granted claim 1 includes a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.
- 2.5.4.1. Granted claim 1 states that in response to a read request, the output drivers of the memory device output a first portion of data synchronously with respect to a rising edge transition of an external clock and a second portion synchronously with respect to the immediately following falling edge transition. PCT claim 103 was silent as to existence of a clock signal and therefore said nothing about the relationship between the clock signal and the timing of data output or the relationship between the clock frequency and the data rate. Granted claim 1, in effect, specifies the use of a system clock running at half the bus cycle rate.
- 2.5.4.2. According to the Grandparent Application, the technical problem to be solved by the use of a system clock running at half the bus cycle rate was to reduce clock distribution problems by, firstly, ensuring all signals on the bus have the same worst case data rate and, secondly, facilitating the labelling of even and odd cycles [Grandparent Application, 48:6-16]. In a typical bus system, where each bus cycle corresponds to one clock cycle, the address, data and control signals on the bus will transition, at worse, once in every clock cycle. The clock signal, by definition, transitions twice in every clock cycle. Thus, as the Grandparent Application explains, if the clock is slowed down to half its typical rate, thereby slowing it down relative to all the other signals on the bus, it will transition no faster than the worst case transition rate of all the other signals on the bus. The further result of slowing down the clock in this way is that all the input receivers of any device connected to the bus must be able to sample address, data or control signals on both edge transitions of the clock [Grandparent Application, 54:9-12].
- 2.5.4.3. As explained by the Grandparent Application, labelling of bus cycles as even or odd is facilitated by the use of a clock running at half the bus cycle rate. However, such labelling is needed only if the devices connected to the bus have to identify whether the current cycle is odd or even. Such identification would be required only if the bus protocol restricts the classes of information that can be transferred over the bus on even or odd clock cycles. In the preferred

embodiment, the restriction is that bus transaction requests can only be initiated on even clock cycles [Grandparent Application, 19:15-21]. The reason for this restriction can only be determined from a detailed analysis of the preferred embodiment.

- 2.5.4.4. In the preferred embodiment, the bus runs at a very high data rate. To give the device input receivers sufficient time to amplify signals received from the bus at this high data rate, two input receivers are used with each input pin [Grandparent Application, 54:8-15; 61:9-13]. The separation of the even and odd cycle input signals is maintained within the device, since the gate delays of the internal logic components do not allow individual components to maintain a throughput equal to the bus cycle rate [Grandparent Application: 19:15-17]. Because of this internal separation and the signal throughput limitations, it is possible to test only one of the two sets of input signals for the beginning of a transaction request. In the preferred embodiment, the even cycle input signals are chosen. This explains why the bus protocol requires the labelling of clock cycles.
- 2.5.4.5. Thus, the justification given in the Grandparent Application for running the system clock at one half the bus cycle rate is predicated on the existence of two sets of input receivers for each device, operating on alternate bus cycles, and on the fact that input signals received by these two sets of input receivers are treated differently within the device.
- 2.5.4.6. Therefore, in accordance with the established case law of the Technical Board of Appeal, restricting the claimed semiconductor device to a device that has a data output rate of twice the clock rate, but not to a device that also has data, address and control input rates of twice the clock rate, amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.
- 2.5.5. Granted claim 1 includes a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.
- 2.5.5.1. PCT claim 103 required a "modifiable access time register... whereby data may be transmitted to said register ... which establishes a predetermined amount of time that said semiconductor device thereafter must wait before using said bus in response to a request."
- 2.5.5.2. Granted claim 8 requires there to be more than one selectable access time, suggesting that this is not a requirement of granted claim 1. However, granted claim 1 does require a programmable access time register. If the access time register is to be programmable, it must apparently be capable of holding more than one value. Moreover, claim 1 apparently does not allow each value that the access time register may hold to result in the same number of clock cycles' delay, because at least one value must be "representative" of its corresponding number of clock cycles. If this is correct, it means that granted claim 1 covers devices in which, if a specific value is stored in the programmable access time register, the device delays a predetermined number of clock cycles and outputs data synchronously with respect to the external clock signal; and in which, if another value is stored, the device does not delay.
- 2.5.5.3. It is plain that the objective to be achieved by requiring programmable access time delays is to avoid collisions on the bus. This objective is only completely achieved if the device always waits a selectable number of clock cycles and outputs data synchronously with respect to the external clock signal. If the effect of granted claim 8 is as discussed above, granted claim 1 is not so limited.
- 2.5.5.4. It would follow that a requirement has been introduced into granted claim 1 that the device need only wait a predetermined number of clock cycles for one of the

possible values to be stored in the access time register. Such introduction would fail the test established by the Technical Board of Appeal. It is inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.5.6. Granted claims 1 and 4 include a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.5.6.1. Granted claim 1 requires two portions of data to be output in response to a read request. The only disclosure in the Grandparent Application of the output of multiple portions of data in response to a read request is in the context of a block transaction request. Claim 4 requires the amount of data output to correspond to block size information received by the memory device.

2.5.6.2. The broadest disclosure of block transaction requests in the Grandparent Application was PCT Claim 38. PCT claim 38 required the block-size value to be included in control information that, together with address information, is bundled into a single request. That is exactly what is described in the Grandparent Application with reference to Fig. 4 [Grandparent Application, 21:21-22:10, esp. 22:9-10; 27:23-30]. No advantage is ascribed in the Grandparent Application to the subject-matter of PCT claim 38, as distinct from the subject-matter of the claims from which it depends. Granted claims 1 and 4 contain no explicit requirement for the block size value to be included with address information in a transaction request.

2.5.6.3. As stated above, no advantage is ascribed in the Parent Application to the subject-matter of PCT claim 38, as distinct from the subject-matter of the claims from which it depends. This is a crucial point. According to the test established by the Technical Board of Appeal for admissibility of intermediate generalisation, it can only be justified if there is an advantage – a solution to a technical problem – disclosed in the Grandparent Application that is completely solved by the intermediate generalisation. If there is no problem and no solution, no advantage disclosed, the generalisation must fail. Such is the case here.

2.5.6.4. In spite of the above, it is not difficult to understand the purpose of PCT claim 38 in the context of a system where high-speed access to blocks of data is desired, and in which all bus transaction requests are six bus cycles long. It would be foolish to require one bus transaction to establish a block size and a second bus transaction to make the block transfer request. This would only increase the block request overhead to which the Grandparent Application refers [Grandparent Application, 15:23-16:7]. This purpose, namely allowing block size selection with no additional overhead, was not disclosed as such, but if it had been, it would not assist the patentee: it is only achieved by integrating the block size value into the transaction request, which explains why the inventors took such pains to allow both very small block sizes (block size 1) and very large block sizes (block size 1024) to be represented by binary encoding using only four bits [Grandparent Application, 28:1-11]. It also explains the use of the word "encodes" in addition to "specifies" in PCT claim 38. Unless the block size value is so integrated, bus overhead will increase.

2.5.6.5. For this reason, if granted claims 1 and 4 do not require that block size information be integrated into a transaction request, then the claims fail the test established by the Technical Board of Appeal. They amount to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.5.6.6. It is worthy of note that it is not enough in the present context to say that the advantage or purpose of PCT claim 38 is to allow block size selection; that is tautologous. It amounts to saying that the advantage of an integer is its own

existence; that the problem to be solved by an integer is its own provision. On that basis, every intermediate generalisation could be justified; but that is not the law.

2.5.7. Granted claim 5 includes a technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.5.7.1. Claim 5 of the Patent requires the clock generation circuitry to include "a delay locked loop coupled to the clock receiver circuitry to generate the internal clock signal." The term "delay locked loop" was never used in the Grandparent Application. It is a wholly new term. No delay locked loop was ever disclosed. The nearest thing to it was probably to be found somewhere in the circuit illustrated in Fig. 12 of the Grandparent Application. The broadest disclosure of that circuit in the Grandparent Application is to be found in claims 78 and 108, but these claims make no mention of anything that could be regarded as a delay locked loop. The broadest disclosure in the Grandparent Application of any sufficiently concrete functional detail of the circuit of Fig. 12 is to be found in claim 79.

2.5.7.2. Clearly, the object to be achieved by the circuit of Fig. 12 was to generate an internal clock signal synchronised to a time half way between the early and late bus clock signals [Grandparent Application, 46:20-47:1; 47:21-48:3]. This objective is only achieved if the features of claim 79 of the Grandparent Application are present. Granted claim 5 is not so limited.

2.5.7.3. For this reason, the introduction into granted claim 5 of a delay locked loop in isolation fails the test established by the Technical Board of Appeal. It amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.6. Having looked at the overall picture by comparing the Patent with the Grandparent Application, it is useful now to look at the detail of when the various impermissible changes in subject-matter took place.

2.7. The Parent Application discloses subject-matter not disclosed in the Grandparent Application as originally filed on two counts, namely impermissible claim broadening and impermissible intermediate claim generalisations.

2.8. Impermissible Claim Broadening

2.8.1. Claim 1 of the Parent Application lacks all the essential elements of the invention as originally disclosed.

2.8.1.1. Claim 1 of the Parent Application is not based upon any of the original claims as filed in the Grandparent Application. This was acknowledged by the patentee in its letter dated September 23, 1999 on the Parent Application. In fact, none of the independent claims of the Grandparent Application as filed shares any substantive subject-matter with claim 1 of the Parent Application. It is impossible to identify the subject-matter of claim 1 of the Parent Application with any of the inventions as originally claimed in the Grandparent Application.

2.8.1.2. Accordingly, whichever independent claim of the Grandparent Application is used as the starting point, that claim has been relieved of its very essence. Into the wholly emasculated claim have been introduced a series of isolated integers including all of the three sub-paragraphs of claim 1 of the Parent Application, none of which relate to the objective originally achieved by any of the claims of the Grandparent Application. Clearly, it is not permissible for the patentee to remove from any independent claim of the Grandparent Application the very integers that allow it to achieve the objective that the inventors had in mind for

it. It amounts to abandoning the claim and reconstructing, *ex post facto*, a new claim including a selection of integers isolated from the original disclosure, whether from the description or the claims, irrespective of their relevance to the original primary objectives of the invention. That is not allowed. Such an amendment must fall all three limbs of the test set out in decision T331/87. It would amount to inadmissible claim broadening contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.9. Impermissible Intermediate Generalisation

2.9.1. Claim 1 of the Parent Application includes a technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.9.1.1. Claim 1 of the Parent Application relates to a semiconductor memory device. As explained in paragraphs 2.6.3.1 – 2.6.3.3 above, the only semiconductor memory devices disclosed in the Grandparent Application are those that are compatible with the highly multiplexed bus and transfer device select information via the bus. Restricting any claim to a semiconductor memory device, but not to the transfer of device select information via the bus, amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.9.2. Claim 1 of the Parent Application includes a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.9.2.1. Claim 1 of the Parent Application states that in response to a read operation, the output drivers of the memory device output data onto an external bus in response to first and second internal clock signals. According to the Grandparent Application, two internal clock signals are generated, one being the complement of the other [Grandparent Application, 58:13-16]. These internal clock signals are used to clock input and output drivers [Grandparent Application, 54:8-12; 55:12-16; 58:21-23; Fig. 10]. In the preferred embodiment, the bus runs at a very high data rate. To give the device input receivers sufficient time to amplify signals received from the bus at this high data rate, two input receivers are used with each input pin [Grandparent Application, 54:8-15; 61:9-13]. The separation of the even and odd cycle input signals is maintained within the device, since the gate delays of the internal logic components do not allow individual components to maintain a throughput equal to the bus cycle rate [Grandparent Application: 19:15-17]. For the same reason, output multiplexers are used [Grandparent Application, 55:12-16].

2.9.2.2. Thus, the justification given in the Grandparent Application for using two internal clock signals is predicated on the existence in each device of two sets of input receivers and a set of output multiplexers, operating on alternate bus cycles.

2.9.2.3. Therefore, in accordance with the established case law of the Technical Board of Appeal, restricting the claimed semiconductor device to a device that has two internal clock signals, but not to a device that also has two sets of input receivers and a set of output multiplexers, amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.9.3. Claims 5, 10 and 11 of the Parent Application, include technical features that have been isolated from their original context and do not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.9.3.1. Claims 5, 10 and 11 of the Parent Application require clock generation circuitry to include a delay locked loop coupled to clock receiver circuitry to generate an

internal clock signal. For the reasons given in paragraphs 2.6.7.1 – 2.6.7.3 above, *mutatis mutandis*, the introduction into a claim of a delay locked loop in isolation amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.9.4. Claims 7 and 9 of the Parent Application, include a technical features that have been isolated from their original context and do not provide a complete solution to any technical problem disclosed in the Grandparent Application.

2.9.4.1. Claims 7 and 9 of the Parent Application require first and second external clock signals. These are shown in the circuit illustrated in Fig. 12 of the Grandparent Application. The broadest disclosure of that circuit and its purpose in the Grandparent Application is to be found in claims 73 and 78, which refer to early and late clock signals.

2.9.4.2. Clearly, the object to be achieved by the use of early and late bus clock signals was to allow an internal clock signal synchronised to a time half way between the early and late bus clock signals to be generated [Grandparent Application, 46:20-47:1; 47:21-48:3]. This objective is only achieved if the features of claim 78 of the Grandparent Application are present. Claims 7 and 9 of the Parent Application are not so limited.

2.9.4.3. For this reason, the introduction into claims 7 and 9 of first and second external clock signals with no further explanation fails the test established by the Technical Board of Appeal. It amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.10. The Patent discloses subject-matter not disclosed in the Parent Application as originally filed on two counts, namely impermissible claim broadening and impermissible intermediate claim generalisations.

2.11. Impermissible Claim Broadening

2.11.1. Granted claim 1 lacks all the essential elements of the invention as disclosed in the Parent Application.

2.11.1.1. Claim 1 of the Patent is not based upon the only independent claim, claim 1, of the Parent Application. In fact, these two claims share no substantive subject-matter.

2.11.1.2. Accordingly, claim 1 of the Parent Application has been relieved of its very essence. Into the wholly emasculated claim have been introduced a series of isolated integers including all of the three sub-paragraphs of granted claim 1. This amounts to abandoning the claim and reconstructing, *ex post facto*, a new claim including a selection of integers isolated from the original disclosure, whether from the description or the claims, irrespective of their relevance to the original primary objectives of the invention. That is not allowed. Such an amendment must fail all three limbs of the test set out in decision T331/87. It would amount to inadmissible claim broadening contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.12. Impermissible Intermediate Generalisation

2.12.1. Granted claim 1 includes a technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.

2.12.1.1. Granted claim 1 covers read/write devices such as RAMs that have programmable access times for read requests only, and not for write requests. For the reasons given in paragraphs 2.6.2.1 – 2.6.2.4 above, *mutatis mutandis*,

this amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.12.2. Granted claim 1 includes a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.

2.12.2.1. Granted claim 1 specifies the use of a system clock running at half the bus cycle rate. For the reasons given in paragraphs 2.6.4.1 – 2.6.4.5 above, *mutatis mutandis*, this amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.12.3. Granted claim 1 includes a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.

2.12.3.1. Granted claim 8 requires there to be more than one selectable access time, suggesting that this is not a requirement of granted claim 1. For the reasons given in paragraphs 2.6.5.1 – 2.6.5.4 above, *mutatis mutandis*, this amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2.12.4. Granted claims 1 and 4 include a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.

2.12.4.1. Granted claim 1 requires two portions of data to be output in response to a read request. The only disclosure in the Parent Application of the output of multiple portions of data in response to a read request is in the context of a block transaction request. Claim 4 requires the amount of data output to correspond to block size information received by the memory device. For the reasons given in paragraphs 2.6.6.1 – 2.6.6.6 above, *mutatis mutandis*, this amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

3. ENTITLEMENT TO DIVISIONAL STATUS

3.1. For precisely the reasons discussed above and because the objectionable added subject-matter was present in the Application as filed (a fundamental error), it follows that the Application and the Patent are not entitled to the benefit of the filing and priority dates of the Grandparent Application or the Parent Application. This conclusion follows from the application of decision T904/97, Headnote Point 1. They take as their filing and priority dates the date on which the Application was actually filed, namely April 26, 2000. No subsequent amendment can affect the position under Art. 76(1) EPC, as indicated in decision T873/94, Reasons, Section I, Para. 4.

3.2. For reasons that are essentially similar to those set out above, and because the objectionable added subject-matter was present in the Parent Application as filed, it follows that the Parent Application is not entitled to the benefit of the filing and priority dates of the Grandparent Application. It takes as its filing and priority dates the date on which it was actually filed, namely September 15, 1999. It follows that the Parent Application was never able to transfer to the Application the benefit of the filing and priority dates of the Grandparent Application.

3.3. As explained above in paragraph 1, the Patent is *prima facie* not entitled to divisional status.

4. LACK OF NOVELTY ART. 54 EPC

- 4.1. Each and every claim of the Patent, having a priority date of April 26, 2000, lacks novelty.
- 4.2. WO91/16680 ("The Grandparent Application"), published October 31, 1991
- 4.2.1. The Grandparent Application discloses every feature legitimately disclosed in the Patent, including all the integers of every claim.

5. LACK OF INVENTIVE STEP ART. 56 EPC

- 5.1. The matter disclosed in the Grandparent Application does not support any claim purporting to cover devices capable of use with buses other than the highly multiplexed bus described or read requests other than read request packets. In proceedings brought in the UK, France and Germany under the patent arising from the Grandparent Application, the patentee has asserted that certain terms in claim 1 of that patent should be afforded an interpretation wider than that supported by the Grandparent Application. The following discussion assumes that a consonant interpretation may be advanced by the patentee in respect of the Patent.
- 5.2. The subject-matter of claims 1-17 is obvious.
- 5.3. Claim 1 of the Patent is a mere collocation of two elements, each known *per se* in semiconductor memory devices, each serving to perform the function for which it is known and each being wholly independent of the others. The two elements in question are a clock signal running at half the data rate and a synchronous device with an access time register. As a matter of law, affirmed in decisions T324/94 and T363/94, because each element is known *per se* and there is no super-additive effect involved in the collocation, it follows that the subject-matter of claim 1 is obvious.
- 5.4. Clock Signal Running at Half the Data Rate
- 5.4.1. The use of a clock signal running at half the data rate is *prima facie* obvious. As discussed in the Patent, one objective achieved by slowing down the clock signal to half the data rate is to reduce the transition rate of the clock signal so that it will transition no faster than the worst case transition rate of all the other signals on the bus. This is an obviously desirable objective. Because slowing down the clock signal is an obviously desirable objective, there can be no invention in slowing down the clock signal. It is a *prima facie* obvious measure.
- 5.4.2. US Patent 4,330,852, Published May 18, 1982 ("Redwine")
- 5.4.2.1. Redwine discloses a serial input/output memory device in which the input/output is performed in units of 256 bits [Redwine, abstract; Fig. 1]. To increase the serial access speed, the memory is organised into two arrays [Redwine, 2:55-60; Fig. 1]. A separate internal clock is derived for each array (F1; F2) from an external clock signal F. The two internal clocks are 180° out of phase with one another [Redwine, 4:62-64; Fig. 3]. Thus, the memory device is able to output serial data via output multiplexer 26 at twice the clock rate of the external clock F [Redwine, 7:17-26]. Redwine increases the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.4.3. "Scalable Coherent Interface", SCI-28Nov88-doc20 ("SCI B")
- 5.4.3.1. As discussed below, SCI is an interface standard used to connect devices, including processors, I/O devices and memory. SCI uses a 16 bit wide synchronous, packetized bus, to carry address, data and control information. The narrow 16-bit data path operates at 2ns per word (500MHz). It is clocked with a

250MHz clock with both clock edges active [SCI B, p. 2, col. 2, 8th complete paragraph]. This scheme helps to keep IC pin count down.

5.5. Synchronous Device with an Access Time Register

5.5.1. Japanese Patent Application Sho 62-71428, Published October 5, 1988, and English Translation ("Yamaguchi")

5.5.1.1. Yamaguchi describes a dual-port RAM that has both random access I/O and serial access I/O capabilities [Yamaguchi, 3:1-3; 7:3-9]. The random access and serial ports IO1 ... IO3, SIO1 ... SIO 3 and address lines A0 ... A1 are provided to allow the device to connect to an external bus. The RAM includes four memory arrays M-ARY1 ... M-ARY4 [Yamaguchi, 8:10-12]. Each memory cell array M-ARY1 ... M-ARY4 comprises $m+1$ word lines and $n+1$ sets of complementary data lines which intersect at $(m+1) \times (n+1)$ memory cells [Yamaguchi, 8:17-20]. The device may be formed on a single chip [Yamaguchi, 7:4-6]. As explained below, the device outputs data synchronously with respect to an edge transition of an external clock signal. Thus, Yamaguchi discloses "A synchronous semiconductor memory device having at least one memory array which includes a plurality of memory cells."

5.5.1.2. The dual-port RAM includes a timing control circuit TC that receives a number of external control signals [Yamaguchi, 18:10-13; Fig. 1]. One such external control signal is an external serial clock signal SC that is generated off-chip [Yamaguchi, 18:13-16]. The external serial clock signal SC is used to ensure stable synchronisation of the serial output operation of the dual-port RAM with the dot rate of a high-resolution, high dot rate external CRT [Yamaguchi, 28:7-15]. The external serial clock signal is therefore a signal of a fixed frequency linked to the dot rate of the external CRT. Thus, Yamaguchi discloses on-chip "clock receiver circuitry for receiving an external clock signal having a fixed frequency."

5.5.1.3. The timing control circuit includes a counter circuit CTR [Yamaguchi, 19:12-15]. A counter is a register that is capable of incrementing or decrementing the value it contains. Many microprocessor internal registers, including program counter (PC) and accumulator (AC) registers are also counters of this kind. An alternative arrangement is discussed in which a register latches the count value and a count-up counter is used to count from zero until its output matches the value in the register [Yamaguchi, 28:19-29:1]. In the preferred embodiment, the counter of Yamaguchi is loaded with a value presented on the parallel I/O lines IO1 ... IO4 on the falling edge of the row address strobe signal RAS\ [Yamaguchi, 21:12-17; 25:8-11]. Once loaded with this value, the counter circuit counts down to zero in synchronism with an internal clock signal, the counter advancing timing signal ϕ_{cp} [Yamaguchi, 21:18-22:3]. This internal clock signal ϕ_{cp} is generated by the timing control circuit TC from, and has the same frequency as, the external serial clock signal SC. It is merely a slightly delayed version of the external serial clock signal SC owing to the gate delays of inverters N1 and N2 and the AND gate AG2 of the timing control circuit TC [Yamaguchi, 23:8-14, Figs. 1 & 3]. The overall result is that the counter circuit CTR counts clock cycles of the external serial clock signal SC [Yamaguchi, 19:12-15].

5.5.1.4. As will be described below, the counter circuit is used to delay the output of serial data by a number of clock cycles corresponding to the value loaded into the counter circuit CTR [Yamaguchi, 27:19-21]. Thus, Yamaguchi discloses an on-chip "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request."

5.5.1.5. In accordance with a further internal clock signal ϕ_c , the serial I/O circuit latches the data presented by data registers DR1 ... DR4 on complementary data lines CDS1 ... CDS4 to the serial data lines SIO1 ... SIO4 and hence to an external bus [Yamaguchi, 19:17-19]. Clock signal ϕ_c is synchronised to external serial clock signal SC [Yamaguchi, 26:17-19]. Therefore, Yamaguchi discloses "a plurality of output drivers for outputting data in response to the read request, wherein data is output synchronously with respect to *(an edge transition of the external clock signal)*." After the first transition of the internal clock signal ϕ_c , it is also used to shift a shift register of a pointer PNT that points to the current position in the data registers DR1 ... DR4, thus accessing the next data in the data registers DR1 ... DR4 [Yamaguchi, 26:19-27:7]. Data is output in accurate synchrony with the external serial clock signal SC only once an internal strobe signal ϕ_{dt} , generated when the counter CTR reaches to zero, has been asserted. Thus, Yamaguchi discloses that the data is output "after the number of clock cycles *(represented by the value stored in the access-time register)* of the external clock signal transpire."

5.5.1.6. As described above, it was obvious on the priority date to clock the memory device of Yamaguchi with a clock signal running at half the CRT dot rate, to reduce the clock signal transition rate.

5.5.1.7. It follows that the subject-matter of claim 1 is obvious.

5.5.1.8. As discussed at length above, Yamaguchi describes a timing control circuit that generates a number of control timing signals, one of which is the internal clock signal ϕ_c . This signal causes the serial I/O circuit SIO to drive new data at each rising edge of ϕ_c . Thus, Yamaguchi discloses that the memory device further includes "clock generation circuitry, coupled to the clock receiver circuitry, to generate an internal clock signal", and that "the output drivers output data in response to the internal clock signal."

5.5.1.9. It follows that the subject-matter of claim 3 is obvious.

5.5.1.10. As explained above, the number of clock cycles delay introduced into a serial read operation by the counter circuit CTR is equal to the value stored in it during the read request. That value can be between 0 and 15. The value stored is therefore "representative of one of a plurality of different delay times."


5.5.1.11. It follows that the subject-matter of claim 8 is obvious.

5.6. US Patent 4,785,428, Published November 15, 1988 ("Bajwa")

5.6.1. Fig. 1 of Bajwa is the functional block diagram of a DRAM controller (consisting of functional blocks 9, 11-13, 15-17) that controls a DRAM array 14. In the preferred embodiment, the DRAM controller itself is a chip, [Bajwa, 2:56], while the DRAM array 14 is a 1Mb conventional DRAM [Bajwa, 4:24]. The combination of the DRAM controller and the DRAM array constitutes a semiconductor memory that interfaces with a synchronous bus 10. Thus, Bajwa discloses "a synchronous semiconductor memory having at least one memory array which includes a plurality of memory cells."

5.6.2. The DRAM controller performs the function of the bus interface of the memory device and is a clocked, self-timed device. The DRAM controller is driven by a clock having a two non-overlapping phase design [Bajwa, 2:52-54]. Four clocks PH1 and PH2 and their inverses PH1I and PH2I are disclosed as controlling the logic function [Bajwa, 2:54-59]. Fig. 4 shows how a memory access sequence is clocked with respect to PH1. Bajwa further describes a clock management circuit 28 and RAS, WE, OE, and CAS pin control logic which receive an external clock in the form of clock signals CLKA, CLKB, CLKC and CLKD [Bajwa, Fig. 2]. The clock

management circuit in combination with the clocking signals controls timing for the memory access [Bajwa, 6:18-23]. Thus, Bajwa discloses "clock receiver circuitry for receiving an external clock signal having a fixed frequency."

- 5.6.3. The memory includes a program RAM 20 for storing "timing sequences programmed into" the DRAM controller to match the speed of any number of existing arrays [Bajwa, 2:44-45; 2:6-7]. The timing sequences are programmed into program RAM 20, which is a 53x35 array as shown in Fig. 2. Each of the 53 rows is used for programming all 35 internal transition control signals required for one clock cycle. The 53 rows are used as follows. The DRAM controller is capable of block size access of 1 to 4 words in size [Bajwa, 8:28-34], and the timing sequence for each of the 4 block size types is programmed separately, i.e. by using:
- 5.6.3.1. 8 rows for the timing sequence of 1-word access, [Bajwa, Fig. 6] (max delay of 8 clock cycles),
 - 5.6.3.2. 11 rows for programming the timing sequence of a 2-word access (max delay of 11 clock cycles),
 - 5.6.3.3. 15 rows for programming the timing sequence of a 3-word access (max delay of 15 clock cycles),
 - 5.6.3.4. 19 rows for programming the timing sequence of a 4-word access (max delay of 19 clock cycles),
- 5.6.4. This requires a total of 53 rows. Thus, program RAM 20 comprises at least four access-time registers used to specify 4 different access times corresponding to each of the 4 possible (programmable) block sizes. In this scheme, one pre-programmed signal RPYNOW signals the memory control unit to transmit the read reply packet to the AP bus [Bajwa, 5:64-6:14]. RPYNOW can correspond to any one of a number of clock cycles [Bajwa, 6:10-15]. Thus, Bajwa discloses a "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request." It further discloses that the value stored is therefore "representative of one of a plurality of different delay times." 
- 5.6.5. The program RAM 20 contains the DRAM access protocols and timings [Bajwa, 8:14-38]. The signal multiplexer 22 selects the appropriate signals from the timing sequence stored in the program RAM 20 and drives the various RAS, CAS, OE, and WE signals, along with the signals from clock management unit 28, clocking signals (CLKA-CLKD) and ARREN [Bajwa, Figs. 1 and 2]. As shown these signals are synchronous with respect to, at least, clock PH1. They are supplied to the DRAM and cause the DRAM to output the selected amount of data. The read data coming in from the DRAM is queued in the memory control unit and is output onto the external bus (AP Bus) according to a pre-programmed timing sequence which indicates the clock cycle during which the transmission of the read data may begin [Bajwa, 5:64-6:14]. The AP bus timing requires the transmission of a data word on every bus cycle [Bajwa, 6:1-2]. Thus, the data is read out synchronously with respect to the external clock signal. Accordingly, Bajwa discloses "a plurality of output drivers for outputting data in response to a read request, wherein the data is output after the number of clock cycles [represented by the value stored in the access-time register] of the external clock transpire and synchronously with respect to [an edge transition of the external clock signal]."
- 5.6.6. It was obvious on the priority date to integrate the timing and interface functions of the MCU onto individual memory devices, to allow higher operating speeds. Moreover, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.

- 5.6.7. It follows that the subject-matter of claim 1 is obvious.
- 5.6.8. When a read request is received, the central sequencing logic (CSL) 17 of the DRAM controller starts a timing sequence that provides all the timings required for accessing the DRAM array in each clock cycle. A SEQUENCE signal is set high to indicate that the DRAM controller is currently processing the request [Bajwa, 3:66-67; Fig. 4]. SEQUENCE is set low (inactive) when the processing of the request is completed. A RAS PRECHARGE is discussed [Bajwa, 7:15-26]. The DRAM RAS precharge time is specified by a 3-bit counter in the preferred embodiment. The counting of the precharge cycles is enabled when SEQUENCE goes inactive [Bajwa, 7:21-22]. Hence, the memory array is "automatically precharged after receiving the read request" by the DRAM controller. When the count expires, precharge is complete and the DRAM controller is ready to process the next request (the precharge counter is also reloaded in preparation for the next read cycle) [Bajwa, 7:23-27].
- 5.6.9. It follows that the subject-matter of claim 2 is obvious.
- 5.6.10. As discussed above, The DRAM controller is a clocked, self-timed device. The DRAM controller is driven by a clock having a two non-overlapping phase design [Bajwa, 2:52-54]. Four clocks PH1 and PH2 and their inverses PH1I and PH2I are disclosed as controlling the logic function [Bajwa, 2:54-59]. Fig. 4 shows how a memory access sequence is clocked with respect to PH1. The signal multiplexer 22 selects the appropriate signals from the timing sequence stored in the program RAM 20 and drives the various RAS, CAS, OE, and WE signals, along with the signals from clock management unit 28, clocking signals (CLKA-CLKD) and ARREN [Bajwa, Figs. 1 and 2]. As shown these signals are synchronous with respect to, at least, clock PH1. They are supplied to the DRAM and cause the DRAM to output the selected amount of data. Thus, Bajwa discloses "clock generation circuitry, coupled to the clock receiver circuitry, to generate an internal clock signal" and that "the plurality of output drivers output data in response to the internal clock signal."
- 5.6.11. It follows that the subject-matter of claim 3 is obvious.
- 5.6.12. The DRAM controller provides timing sequences for one to four word memory accesses [Bajwa, 8:16-19]. During a memory operation, the DRAM sequence controller 24 receives a two bit block size information code (NUMWORDS) which corresponds to the number of words for the memory access (1 to 4 words) [Bajwa, 3:63-66; Fig. 2]. A 3-bit unary encoded or 2-bit binary encoded operation code signal is supplied to a SIGNAL MUX 22 to specify whether the operation is a read, write or refresh [Bajwa, 4:6-10]. Thus, Bajwa discloses that "the data output in response to the read request corresponds to an amount of data specified by block size information, wherein the block size information is provided to the memory device."
- 5.6.13. It follows that the subject-matter of claim 4 is obvious.
- 5.6.14. As explained above, The DRAM controller is capable of block size access of 1 to 4 words in size [Bajwa, 8:28-34], and the timing sequence for each of the 4 block size types is programmed separately. The value stored in the access time register is therefore "representative of one of a plurality of different delay times."
- 5.6.15. It follows that the subject-matter of claim 8 is obvious.
- 5.7. GigaBit Logic, 1988 GaAs IC Data Book & Designer's Guide, May 1988, 12G014 256x4-Bit Registered, Self-Timed Static RAM ("GigaBit")
- 5.7.1. Gigabit describes a 256x4-bit static RAM fabricated using Gallium Arsenide (GaAs) technology ("The 12G014"). The 12G014 is a self-timed SRAM ("STRAM"). Thus,

Gigabit discloses "A synchronous semiconductor memory device having at least one memory array which includes a plurality of memory cells."

- 5.7.2. The 12G014 is a self-timed SRAM ("STRAM"), running at a fixed frequency of 400 MHz. The clock cycle is therefore 2.5 ns. The 12G014 STRAM has differential clock inputs CLK and CLK \bar , as shown in the block diagram on page 2-3. The differential clock inputs connect to an output clock generator functional block, which must include clock receiver circuitry to receive the differential clock inputs. Thus, Gigabit discloses "clock receiver circuitry for receiving an external clock signal having a fixed frequency."
- 5.7.3. The 12G014 is a $\times 4$ memory, with four output drivers (Q0-Q3) outputting data onto the bus in response to a read request. The output is fully registered (double latch) as shown in the block diagram on page 2-3. The output register is clocked synchronously to the external clock via the agency of the output clock generator that generates an internal clock from the complementary clock inputs CLK and CLK \bar . Hence, the output operation is done synchronously with respect to the external clock. Thus, Gigabit discloses a device that includes "a plurality of output drivers for outputting data in response to a read request ... and synchronously with respect to [an edge transition of the external clock signal]."
- 5.7.4. The 12G014 has three output modes, namely latch mode, register mode and transparent mode. The transparent mode is asynchronous (similar to conventional SRAMs). Both the latched mode and register mode are clocked. In register mode, memory access takes place during the clock cycle in which the read request is received. Valid output data is presented to the on-chip output drivers. At the next rising edge of the clock signal, i.e. a full 2.5 ns clock cycle after the read request is received, this data is loaded into the output drivers and propagates to the data output pins, and hence the bus, where it is held for a full cycle. This is shown and described in the timing diagram on page 2-6. Thus, Gigabit discloses a device in which, in register mode, the output drivers output data "after [a] number of clock cycles [i.e. one] of the external clock transpire and synchronously with respect to [an edge transition of the external clock signal]."
- 5.7.5. In latch mode, memory access again takes place during the clock cycle in which the read request is received. Valid output data is presented to the on-chip output drivers. However, these output drivers are driven transparent at the falling edge of the clock signal to allow valid data to appear on the data output pins, and hence on the bus, as soon as possible. The duty cycle of the clock is changed so that the falling edge arrives after less than one half clock cycle. The output drivers are driven transparent at this falling edge and latched at the next rising edge to hold data over to the falling edge of the next clock cycle. This is shown and described in the timing diagram on page 2-7. Thus, Gigabit discloses a device in which, in latch mode, the output drivers output data "after [a] number of clock cycles [i.e. less than one half] of the external clock transpire and synchronously with respect to [an edge transition of the external clock signal]."
- 5.7.6. The output mode is programmed by applying one of three signal levels to a MODE pin, as described on page 2-4. Vss level gives register mode, Vdd gives latch mode and Vcc gives transparent mode. The mode pin signal is received in the output clock generator where it must be decoded. The circuitry that decodes the mode pin signal and produces the decoded logical outputs is programmable until the mode pin is connected and thereafter outputs the value programmed into it. The value programmed into this register determines the output mode and, in particular, the output delay. Therefore, Gigabit discloses an on-chip "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request." It also discloses that, in both latch and register modes, the output drivers output data "after the number of clock cycles of the external clock

transpire and synchronously with respect to *an edge transition of the external clock signal*."

- 5.7.7. As described above, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate; or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.7.8. It follows that the subject-matter of claim 1 is obvious.
- 5.7.9. The Fig. on page 2-3 includes an "Output Clock Generator," which generates internal clock signals from CLK and CLK\ These internal clock signals drive the output register. Thus, Gigabit discloses that the memory device further includes "clock generation circuitry, coupled to the clock receiver circuitry, to generate an internal clock signal" and that the "output drivers output data in response to the internal clock signal."
- 5.7.10. It follows that the subject-matter of claim 3 is obvious.
- 5.7.11. As discussed above, the value programmed into the access-time register determines the output mode and, in particular, the output delay. It is representative of one of a plurality of different delay times."
- 5.7.12. It follows that the subject-matter of claim 8 is obvious.
- 5.7.13. As illustrated in the block diagram on page 2-2, all signal inputs to the device are latched in accordance with an internal clock generated by an input clock generator from the differential external clock signals CLK and CLK\ Thus, Gigabit discloses that the "read request is sampled synchronously with respect to a rising edge transition of the external clock signal."
- 5.7.14. It follows that the subject-matter of claim 10 is obvious.
- 5.7.15. The Pin descriptions of page 2-4 of GigaBit describe pin VBB as an "ECL Threshold Reference Voltage (Nominal - 1.3V) provided to allow direct tracking of an ECL logic family's Vbb reference voltage." ECL signals are low voltage swing signals. GigaBit continues: "This must be tied to a nominal - 1.3V when interfacing to IDG PicoLogic circuitry (or to the VBBS supply pin of a PicoLogic device)." CLK\ is described thus: "Clock Input Complement: provides an optional differential input for systems with differential Clock distribution. For single ended Clock distribution CLK\ is tied to the Vbb input pin (nom. - 1.3V). Page 2-5 mentions that CLK\ can be used instead of CLK, in which case CLK is tied to a nom - 1.3V source. It is clear from these passages that both CLK and CLK\ are low voltage swing clock signals.
- 5.7.16. It follows that the subject-matter of claim 17 is obvious.

 5.8. US Patent 4,858,113, Issued August 15, 1989, ("Saccardi")

- 5.8.1. Saccardi describes a massively parallel processor architecture in which a plurality of components are interconnected by a cross bar switch 120 [Saccardi, Fig. 4]. Some of these components are semiconductor memory devices 108, 110. These memory devices are addressed by other devices via the cross bar [Saccardi, 4:51-53] and output or input word-size read data via the cross bar to or from other devices [Saccardi, 4:27-28; 4:54-59]. Independent read and write requests are possible. The inter-device connections and connections between the cross bar and the memory devices form a bus external to the memory devices. Each device therefore includes at least one memory array containing a plurality of memory cells. As explained below, the device outputs data synchronously with respect to an edge transition of an external clock signal. Thus, Saccardi discloses "A synchronous semiconductor

memory device having at least one memory array which includes a plurality of memory cells."

- 5.8.2. Each memory device includes a variable tick delay device VTD [Saccardi, 3:25:27]. Separate simplified devices are an alternative [Saccardi, 3:61:63]. The VTDs control the input and output operations of the memories, introducing variable delays to align data for processing in a selected sequence relative to the clocking period [Saccardi, 3:26-35]. All the memory devices and arithmetic units operate under timing control of a single clock. The clocking rate is predetermined [Saccardi, 3:64-7:3]. Thus, Saccardi discloses "clock receiver circuitry for receiving an external clock signal having a fixed frequency."
- 5.8.3. Fig. 6 of Saccardi is a block diagram of a VTD. As shown, there is a register comprising 32 bits of control, C0-C31. Each group of 4 bits, e.g. C0-C3, defines an access delay that can be 1-16 clock cycles long. The VTD that is at the output of a memory device thus delays the memory response by a programmed number of clock cycles [Saccardi, 4:5-8]. Thus, Saccardi discloses a "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request."
- 5.8.4. Fig. 6 also shows a 64-bit wide data path, hence there is a plurality of output drivers. These are shown as output registers 17 [Saccardi, 3:45-50]. A system clock provides timing control for data transfer and for functional operation of each element in the system and this has to include the output registers of the VTD [Saccardi, 3:22-25]. Hence, data is output by a VTD in a delayed manner, after a programmable number of clock cycles, and synchronously with respect to the external system clock signal. Thus, Saccardi discloses "a plurality of output drivers for outputting data in response to a read request, wherein the data is output after the number of clock cycles *represented by the value stored in the access-time register* of the external clock transpire and synchronously with respect to *an edge transition of the external clock signal*."
- 5.8.5. As described above, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.8.6. It follows that the subject-matter of claim 1 is obvious.
- 5.8.7. In the pipelined processor model, the delays of the VTDs are established in advance and the pipeline is then clocked through. Setting up the control registers of the VTDs amounts to an "initialisation sequence" and clearly must be performed "after power is applied." Moreover, the transmission to the VTD of control register set-up data amounts to making a "set register request" that is distinct from subsequent read requests.
- 5.8.8. It follows that the subject-matter of claims 6 and 7 is obvious.
- 5.8.9. As explained above, the number of clock cycle delays introduced into a read operation by the VTD is one of a number of possible delays depending on the value stored in the control register during the set register request. In Fig. 6, that value can be between 0 and 15. The value stored is therefore "representative of one of a plurality of different delay times."
- 5.8.10. It follows that the subject-matter of claim 8 is obvious.

5.9. Japanese Patent Application S57-14992, Published January 26, 1982, and English Translation ("Taguri")

- 5.9.1. Fig. 2 of Taguri is the functional block diagram of a memory device MS 1 [Taguri, 2:12-14]. The MS1 includes an input latch 2, an output latch 3, a controller 4 and a memory array 5 [Taguri, 3:22-23]. Interface signals S_i ($i=1-n$) sent from a CPU 6 are latched to the input latch 2 by clock signal t_i [Taguri, 3:23-24]. The output latch 3 sends data to the CPU 6 as m interface signals S_o ($o=1-m$) by clock signal t_j [Taguri 3:26-27]. Thus, Taguri discloses "a synchronous semiconductor memory device having at least one memory array which includes a plurality of memory cells."
- 5.9.2. Clock signals t_i and t_j are selected from four external clocks t_0-t_3 , phase shifted relative to one another in quarter cycle increments [Taguri, 3:1-5; 3:28-29]. Thus, Taguri discloses "clock receiver circuitry for receiving an external clock signal having a fixed frequency."
- 5.9.3. The selection of the appropriate clocks t_i and t_j from the four external clocks t_0-t_3 is done by a clock selector 8, controlled by a configuration control register 7 [Taguri, 3:30-31]. The configuration control register 7 is programmable [Taguri, 3:34-37]. This allows the time delay between receipt of a read request and data output to be set at will in quarter cycle increments [Taguri, 4:1-3]. Thus, Taguri discloses a "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request."
- 5.9.4. As discussed above, the output latch 3 is clocked by clock signal t_j . The output is m bits wide. Thus, Taguri discloses "a plurality of output drivers for outputting data in response to a read request, wherein the data is output after the number of clock cycles [represented by the value stored in the access-time register] of the external clock transpire" and "synchronously with respect to [an edge transition of the external clock signal]."
- 5.9.5. It was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.9.6. It follows that the subject-matter of claim 1 is obvious.
- 5.9.7. The clock selector circuit 8 selects one of the four external clocks to control output and one to control input. Thus, Taguri discloses "clock generation circuitry, coupled to the clock receiver circuitry, to generate an internal clock signal" and that "the plurality of output drivers output data in response to the internal clock signal."
- 5.9.8. It follows that the subject-matter of claim 3 is obvious.
- 5.9.9. The time delay between receipt of a read request and data output can be set at will in quarter cycle increments [Taguri, 4:1-3]. Thus, Taguri discloses that the value stored in the programmable access-time register is "representative of one of a plurality of different delay times."
- 5.9.10. It follows that the subject-matter of claim 8 is obvious.
- 5.9.11. Interface signals S_i ($i=1-n$) sent from a CPU 6 are latched to the input latch 2 by clock signal t_i [Taguri, 3:23-24]. Thus, Taguri discloses that the "read request is sampled synchronously with respect to a rising edge transition of the external clock signal."
- 5.9.12. It follows that the subject-matter of claim 10 is obvious.

5.10. US Patent No. 4,499,536, Issued February 12, 1985 ("Gemma")

- 5.10.1. Gemma describes a processor-based SCU that interfaces with a main memory including a plurality of memory cells arranged as one or more memory arrays [Gemma, 3:65]. The memory arrays may be in a number of configurations [Gemma, 6:25-32]. The SCU includes a controller 17 that receives a clock signal T_2 synchronised with the processor machine cycle [Gemma, 4:26-35; Fig. 2]. This is received by AND gates 29 and 24 within the controller 17. Timing control signals that drive a counter 25 are generated from the clock signal T_2 [Gemma, Fig. 2]. As explained below, the device outputs data synchronously with respect to an edge transition of the clock signal T_2 . Thus, Gemma discloses "a synchronous semiconductor memory having at least one memory array which includes a plurality of memory cells" that includes "clock receiver circuitry for receiving an external *(to the SCU)* clock signal having a fixed frequency."
- 5.10.2. The SCU 17 receives memory related instructions from a current instruction register via a signal line 101 under the control of an execution unit. Control information is sent to the SCU from the execution unit via a signal line 102. The control information determines the "type of main memory access, that is, read (FE), full write (ST) and partial write (PST)" [Gemma, 3:10-17]. The generation of a main memory access start signal (EX) that is sent to the main memory is described. The signal has a pulse width predetermined by the timing of timing signals T_0 and T_1 [Gemma, 3:18-41]. Subsequently, the SCU 19 sends a GFDR signal to the main memory on signal line 19 to instruct the main memory to send readout data to a data bus 118 and sends an ADV signal to the processor on signal line 120 to instruct a read data register in the processor to read the data on the data bus 118 [Gemma, 3:52-56]. The content of the read data register is sent to the execution unit 2 via a data bus 122 [Gemma, 3:58-60]. It is implicit in the requirement for the main memory to respond to the GFDR signal that the main memory includes output drivers that respond to the signal. These output drivers output data onto the data bus 118. Thus, Gemma discloses "a plurality of output drivers for outputting data onto an external *(to the main memory)* bus in response to a read request."
- 5.10.3. The SCU includes a controller 13, which is shown in detail in Fig. 2. The controller 13 has "configuration registers 20 and 21 which retain identification flags for the machine cycle of the processor and identification flags concerning the access time of the main memory cells, respectively" [Gemma, 3:61-65]. As described, the controller includes "a counter 25 for counting an elapsed time after the signal EX of a predetermined pulse width has been produced, a decoder 26 for decoding the count of the counter 25, a control circuit 27 for producing the signals GFDR and ADV ... based on the output of the decoder 26" [Gemma, 4:4-10]. Therefore, the timing of the GFDR and ADV signals is determined with reference to the start of the count of the counter. The count is begun with reference to the EX signal that has a predetermined pulse width [Gemma, 3:39-41]. The counter counts according to the clock signal T_2 synchronised with the machine cycles [Gemma, 4:26-35]. Thus, in responding to the GFDR signal, the output drivers of the main memory output data "after *[a]* number of clock cycles of the external *(to the SCU and memory)* clock transpire and synchronously with respect to *[an edge transition of the external clock signal]*."
- 5.10.4. The configuration registers 20 and 21 of the controller 13 are initialised to contain one of three machine cycle identification flags and one of three main memory identification flags [Gemma, 4:11-17]. The output from the configuration registers 20 and 21 are supplied to the control circuit 22 via signal lines 200 and 201 [Gemma, 4:22-25]. The counter 25 is reset in response to the memory signal EX being low [Gemma, 4:32-35]. The output of the counter, representing a number in binary form, 20-23, is fed to a decoder 26 [Gemma, 4:35-38].
- 5.10.5. According to Gemma, "the relation between the combinations of the identification flags m_k ($k=1\sim3$) and the t_{aj} ($j=1\sim3$) and the send timings of the signals GFDR, ADV and BSYR. The send timing are represented by the counts $C_1\sim C_{16}$ of the number of

machine cycle counted after the signal EX has been produced. Those counts are predetermined based on the response performance of the processor and the main memory, i.e. based on the contents of the configuration registers [Gemma, 4:53-60]. The following example of the timing is given: "when the identification flags m_1 and ta_1 are set to "1", respectively, the signal GFDR is sent at the timing C_i and the signal ADV is sent at the timing C_{i+1} " [Gemma, 4:60-63]. Thus, once the EX signal has been produced, that is, once the memory has been instructed to execute a read, the signal GFDR is sent to the memory to instruct it to output the result of the read at a time of C_i which is measured in terms of machine cycles of the processor, the number being determined by the contents of the configuration registers. Having issued the signal GFDR, the controller then issues the signal ADV at a time of C_{i+1} to instruct the read data register to read the output from the memory. Accordingly, Gemma discloses more than one "programmable access-time register for storing a value which is representative of a number of clock cycles of the external [to the SCU] clock signal to transpire after which the memory device responds to a read request." It also discloses that the output drivers of the main memory output data "after the number of clock cycles of the external [to the SCU and the memory] clock transpire and synchronously with respect to [an edge transition of the external clock signal]."

- 5.10.6. It was obvious on the priority date to integrate the timing and interface functions of the SCU onto individual memory devices, to allow higher operating speeds. Moreover, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.10.7. It follows that the subject-matter of claim 1 is obvious.
- 5.10.8. As discussed extensively above, AND gate 24 produces an internal clock signal that is used to drive the counter 25 on a rising edge transition. Thus, Gemma discloses "clock generation circuitry, coupled to the clock receiver circuitry, to generate an internal clock signal", and that "the output drivers output data in response to the internal clock signal."
- 5.10.9. It follows that the subject-matter of claim 3 is obvious.
- 5.10.10. The initialisation of the configuration registers 20 and 21 is described as follows. In the initialisation of the processor, one of the identification flags $m_1 \sim m_2$ and one of the identification flags $ta_1 \sim ta_2$ are set to "1". The operation may be carried out by a known technique, such as by loading a microprogram into the control memory during the initialisation of the processor. [Gemma, 4:9-22]. Thus, Gemma discloses that "the value is stored in the programmable access time register after power is applied to the memory device during an initialisation sequence." This process amounts to a "set register request" from the processor.
- 5.10.11. It follows that the subject-matter of claims 6 and 7 is obvious.
- 5.10.12. The values, $m_1 \sim m_2$ and $ta_1 \sim ta_2$, stored in the configuration control registers clearly control the timing of the output signals. Thus, Gemma discloses that "the value stored in the programmable access time register is representative of one of a plurality of different delay times."
- 5.10.13. It follows that the subject-matter of claim 8 is obvious.



5.11. Japanese Patent Application Sho 62-185253, Published January 31, 1989, and English Translation ("Kumagai")

- 5.11.1. Kumagai discloses a main storage unit MS 4 that includes RAM memory arrays RAM0, RAM1, RAM2, RAM3 [Kumagai, Fig. 1]. The memory device MS 4 is

clocked by an external clock source 5. That clock source is at least common to the memory device MS 4 and the memory controller SCU. It is a fixed frequency clock [Kumagai, Fig. 4]. The clock is used to clock all the interfaces of MS 4 which are all latched: command/address buffer MRQ 20, input data MSD 21 and output data MFD 30 [Kumagai, 4:8-22]. Thus, Kumagai discloses "a synchronous semiconductor memory having at least one memory array which includes a plurality of memory cells."

- 5.11.2. The circuitry to receive the external clock is not specified in Kumagai, but it must exist and can simply be an input buffer. Thus, Kumagai discloses "clock receiver circuitry for receiving an external clock signal having a fixed frequency."
- 5.11.3. The control unit of memory device MS 4, MCR 51, contains the clock counter circuitry shown in Fig. 3. The circuitry includes latches C0 ... C3. These latches store a value representative of a number of clock cycles of the external clock, which is input to the circuitry at 316 (clock T0). The combination of blocks 300, 301, 318 and 319 of Fig. 3 is a clock counter that would signal (see input to block 102) when the clock count reaches the delay value stored in latches C0-C3. That signal then is used to trigger CAS₁ (through the clock phase selection circuitry 102, 303, 304 and 305). Hence, data in RAM0-3 is accessed only after the number of clock cycles stored in C0-C3 has transpired. The data is then sent to the memory output latch MFD 30 to be output at a clock edge to the requesting device SCU 3 [Kumagai, 5:10-27]. Latches C0-C3 are programmable in the sense that they hold whatever values were read into them. Because C0-C3 are part of MCR 51, which is programmed by SCU 3 via MRQ 20, C0-C3 are programmable by the SCU 3 via that interface. Thus, Kumagai discloses a "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request."
- 5.11.4. Each of the synchronous interfaces to the memory device MS 4 of Kumagai, namely the command/address interface at MRQ 20, the input data interface at MSD 21 and the output data interface MFD 30 can each be one or many bits wide. Thus, Kumagai discloses at least one "output driver for outputting data." That output operation is in response to a read request from SCU 3 via command interface MRQ 20 [Kumagai, 4:32-5:2]. The output operation is delayed until after a specified number of clock cycles has transpired. The output operation is synchronous with respect to the external clock because MFD 30 is a latch. Thus, Kumagai discloses that the one or more output drivers of the main memory output data "after the number of clock cycles of the external clock transpire and synchronously with respect to *[an edge transition of the external clock signal]*" [cf. Kumagai, 5:27].
- 5.11.5. It was obvious on the priority date to integrate the timing and interface functions of the SCU onto individual memory devices, to allow higher operating speeds. Moreover, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.11.6. It follows that the subject-matter of claim 1 is obvious.
- 5.11.7. Kumagai discloses an MS control unit MCR that generates timing signal 65 used to latch the fetch data latches FDR1 ... FDR4 [Kumagai, 4:37-39]. This timing signal 65 is an internal clock signal and, in the multi-bit system discussed above where the fetch data latches correspond to the output drivers of granted claim 1, it would constitute an "internal clock signal" generated by "clock generation circuitry, coupled to the clock receiver circuitry." The "output drivers" would "output data in response to the internal clock signal."
- 5.11.8. It follows that the subject-matter of claim 3 is obvious.

5.11.9. The value for access time in Kumagai is stored in latches C0-C3. Latches are devices that hold values clocked in after power is applied. The latches C0-C3 are programmed via the command interface MRQ 20. In most cases, it would be done once and for all after power up, because in a given system the memory access time and processor machine cycle will not change. This amounts to programming "during an initialization sequence of the memory device following power up."

5.11.10. It follows that the subject-matter of claim 6 is obvious.

5.11.11. The latches C0-C3 of Kumagai allow four different access times to be programmed into the memory device MS 4 [Kumagai, 5:27, Fig. 3: blocks 300, 318, 319 and 314].

5.11.12. It follows that the subject-matter of claim 8 is obvious.

5.11.13. As discussed above, the common clock is used to clock all the interfaces of MS 4 which are all latched: command/address buffer MRQ 20, input data MSD 21 and output data MFD 30 [Kumagai, 4:8-22]. Thus, Kumagai discloses that the "read request is sampled synchronously with respect to a rising edge transition of the external clock signal."

5.11.14. It follows that the subject-matter of claim 10 is obvious.



5.12. Japanese Patent Application 60-80193, Published May 8, 1985, and English Translation (Hasegawa)

5.12.1. Hasegawa discloses a memory system having a plurality of storage device blocks. These may be fast RAMs, slow RAMs or ROMs. RAMs and ROMs are designed in array form and equipped with memory cells. As explained below, the device outputs data synchronously with respect to an edge transition of an external clock signal. Thus, Hasegawa discloses "a synchronous semiconductor memory having at least one memory array which includes a plurality of memory cells."

5.12.2. A shift register 5 receives an external clock signal ϕ [Hasegawa, Fig. 1; 4:7-11]. This clock signal ϕ has a fixed frequency [Hasegawa, Fig. 2]. Thus, Hasegawa discloses "clock receiver circuitry for receiving an external clock signal having a fixed frequency."

5.12.3. Registers 2 and 3 hold different values defining how many clock cycles of the clock signal ϕ must elapse after a read request, before data is output via the data bus DB. This is described for a slow memory M2 [Hasegawa, Fig. 2; 4:5-29]. It is also described for a fast memory [Hasegawa, Fig. 3; 4:30-5:10]. Depending upon which memory is accessed, the value stored in either the register 2 or the register 3 is loaded into the shift register 5. The contents of shift register 5 are then right shifted until a zero and a one are present at the outputs D7 and D8 of the shift register 5, whereupon an AND gate G4 asserts $\phi 2$ [Hasegawa, 3:35-39]. Hasegawa also suggests the use of a programmable counter instead [Hasegawa, 5:36-38]. Thus, Hasegawa discloses "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request."

5.12.4. Output drivers are present in the latch register 8 and buffer 9, to output data onto the data bus DB. The clock signal $\phi 2$ is generated synchronously with the external clock signal ϕ [Hasegawa, Figs. 2 and 3]. Thus, Hasegawa discloses "a plurality of output drivers for outputting data in response to a read request, wherein the data is output after the number of clock cycles [represented by the value stored in the access-time register] of the external clock transpire" and "synchronously with respect to [an edge transition of the external clock signal]."

- 5.12.5. It was obvious on the priority date to integrate the timing control functions onto individual memory devices, to allow higher operating speeds. Moreover, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, and to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.12.6. It follows that the subject-matter of claim 1 is obvious.
- 5.12.7. As discussed above, the clock signal $\phi 2$ is generated synchronously with the rising edge of the external clock signal ϕ . The output drivers are driven by $\phi 2$. Thus, Hasegawa discloses "clock generation circuitry, coupled to the clock receiver circuitry, to generate an internal clock signal" and that "the plurality of output drivers output data in response to the internal clock signal."
- 5.12.8. It follows that the subject-matter of claim 3 is obvious.
- 5.12.9. As explained above, registers 2 and 3 can hold a number of different values. Thus, the value stored in the access time register is "representative of one of a plurality of different delay times."
- 5.12.10. It follows that the subject-matter of claim 8 is obvious.
- 5.13. Japanese Patent Application No. S56-82961, Published July 7, 1981, and English Translation ("Kawamasa").
- 5.13.1. Kawamasa discloses a memory system having two memory blocks of differing access times. The memory blocks are designed in array form and equipped with memory cells [Kawamasa, Fig.]. As explained below, the device outputs data synchronously with respect to an edge transition of an external clock signal. Thus, Kawamasa discloses "a synchronous semiconductor memory having at least one memory array which includes a plurality of memory cells."
- 5.13.2. A memory access cycle counter 7 receives an external clock [Kawamasa, 4:10-12]. Thus, Hasegawa discloses "clock receiver circuitry for receiving an external clock signal having a fixed frequency."
- 5.13.3. At the start of each read request, counter 7 is loaded with the access time for the memory being accessed, and then it is counted down at each clock cycle of the external clock. The requested data is output synchronously onto an external bus when the count reaches zero. This is achieved by "0" detector 9 activating AND gate 5, which produces a control signal for latching the data output from the memory into output register 6 [Kawamasa, 4:16-20]. Whenever the CPU accesses a "type 1" high speed memory, counter 7 is programmed with a "fixed" value, which is itself stored in another register 11, that is representative of the access speed of the high speed memory [Kawamasa, p. 3]. If the CPU accesses a slow speed "type 11" memory, a different value is computed and loaded into the counter 7 [Kawamasa, p. 4]. Thus, Kawamasa discloses "a programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request."
- 5.13.4. The "fixed" access time value stored in register 11, corresponding to the access time of the high speed memory "type 1", can itself be programmed using the service processor 16 [Kawamasa, p. 5]. Register 11 also therefore constitutes "a programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory device responds to a read request." This programmable access time register is specific to the high speed memory.

- 5.13.5. Output drivers are present in the output register 6, to output data. The control signal for the output register is generated synchronously with the external clock. Thus, Kawamasa discloses "a plurality of output drivers for outputting data in response to a read request, wherein the data is output after the number of clock cycles (represented by the value stored in the access-time register) of the external clock transpire" and "synchronously with respect to (an edge transition of the external clock signal)."
- 5.13.6. It was obvious on the priority date to integrate the timing control functions onto individual memory devices, to allow higher operating speeds. Moreover, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.13.7. It follows that the subject-matter of claim 1 is obvious.
- 5.13.8. As discussed above, the output of AND gate 5 is generated synchronously with the external clock signal. Thus, Kawamasa discloses "clock generation circuitry, coupled to the clock receiver circuitry, to generate an internal clock signal" and that "the plurality of output drivers output data in response to the internal clock signal."
- 5.13.9. It follows that the subject-matter of claim 3 is obvious.
- 5.13.10. As explained above, registers 7 and 11 can hold a number of different values. Thus, the value stored in the access time register is "representative of one of a plurality of different delay times."
- 5.13.11. It follows that the subject-matter of claim 8 is obvious.
- 5.14. B. Ramakrishna Rau et al., "Cydra 5 Departmental Supercomputer Design Philosophies, Decisions, and Trade-offs" Computer IEEE, Jan. 1989, pp. 12-35 ("Cydra 5")
- 5.14.1. Cydra 5 discusses in detail the Cydra 5 departmental supercomputer. The performance of the supercomputer heavily depends on the memory architecture [Cydra 5, "The main memory system", 26:col. 3 - 30:col. 3]. Where non-interleaved memory does not provide the necessary bandwidth, two alternative memory architectures are discussed. For memory accesses where the addresses form a contiguous sequence, sequentially interleaved memory works very well [Cydra 5, "Sequentially interleaved memory architectures", 29:col. 2 - 30:col. 1]. For memory accesses with a characteristic stride, such as accesses to data arrays, the addresses might be distributed in such a way that many succeeding memory accesses hit the same memory module of a sequentially interleaved memory architecture. In these circumstances, sequential interleaving does not improve memory bandwidth [Cydra 5, "The stride problem", 30:col. 1]. To deal with this situation, a pseudorandomly interleaved memory architecture is discussed [Cydra 5, "Pseudorandomly interleaved memory architecture", 30:col. 1 - col. 3]. In this architecture, high bandwidth is ensured, but the trade-off is that where the high bandwidth is utilised, memory access times are less predictable and may increase [Cydra 5, 30:col. 2, last paragraph et seq.]. This is due to multiple requests being outstanding against the same memory module, in spite of the pseudorandomly interleaved architecture, where the bandwidth and hence the request rate are high. Of course, as Cydra 5 discusses, the same situation arises in a sequentially interleaved memory architecture, where the memory accesses have a characteristic stride, or in a non-interleaved memory architecture where the memory accesses are sequential. As explained below, the memory system outputs data synchronously with respect to an edge transition of an external clock signal. Thus, Cydra 5 discloses "A synchronous semiconductor

memory having at least one memory array which includes a plurality of memory cells."

- 5.14.2. To deal with the increase and unpredictability in memory access times, Cydra 5 proposes to buffer the requests on busy modules so that the processor need not wait until a busy module becomes free [Cydra 5, 32: Fig. 9]. The problem remaining is how to forecast when the data will be ready for the processor. This is a crucial issue since the compiler must schedule the processor tasks, at compile time, and must therefore make assumptions about how long it will take for data to be ready following a memory access. One scheme would be to assume a fixed latency. However, if the fixed latency is too short, the processor will still need to wait, and if it is too long, the access times will be unnecessarily dilated. Simulations have shown that the optimum memory latency value is very sensitive to the request rate [Cydra 5, "The memory latency register", 30:col. 3]. Therefore, Cydra 5 advocates an alternative scheme using a memory latency register (MLR). The MLR is a programmable register that holds the value of memory latency assumed by the compiler when scheduling the currently executing code. The memory system delays data output by as many external clock cycles as are indicated by the value stored in the MLR [Cydra 5, 33:col. 1]. It follows that the memory must include "clock receiver circuitry for receiving an external clock signal having a fixed frequency."
- 5.14.3. The MLR allows the compiler to treat memory accesses as having a deterministic latency, but to use different values for the latency in different portions of the code so as always to deliver near-optimal performance [Cydra 5, 33:col. 1]. From this, it can be inferred that the code generated by the compiler must include, as appropriate, instructions directing the microprocessor to reprogram the MLR when there is a change in the character or bandwidth of memory accesses. The MLR is plainly a "programmable access-time register for storing a value which is representative of a number of clock cycles of the external clock signal to transpire after which the memory responds to a read request." Moreover, since any memory must be furnished with output drivers of some description, Cydra 5 discloses "a plurality of output drivers for outputting data in response to the read request."
- 5.14.4. As discussed above, "data is output synchronously with respect to *[an edge transition of the external clock signal]*" and "after the number of clock cycles *[represented by the value stored in the access-time register]* of the external clock signal transpire."
- 5.14.5. Cydra 5 does not disclose the location of the MLR, but in any case, it was obvious on the priority date to integrate the memory latency register onto individual memory devices, to allow higher operating speeds as is taught by Yamaguchi. Moreover, it was obvious on the priority date to clock any synchronous memory device with a clock signal running at half the bus data rate, to reduce the clock signal transition rate, or to increase the data output rate to twice the clock speed, in the interests of high-speed data transfer.
- 5.14.6. It follows that the subject-matter of claim 1 is obvious.
- 5.14.7. As explained above, the number of clock cycles delay introduced into a read operation by the MLR is equal to the value stored in it. Values of 17 and 26 cycles are suggested [Cydra 5, 33:col. 1]. The value stored is therefore "representative of one of a plurality of different delay times."
- 5.14.8. It follows that the subject-matter of claim 8 is obvious.

5.15. Other Relevant Documents

5.15.1. Programmable Access Time Delay (Claim 1)

5.15.1.1. US Patent No. 4,445,204, Issued April 24, 1984, ("Nishiguchi")

** No teaching of a programmable register storing a binary value in response to a first operation code.*

5.15.1.1.1. Nishiguchi is directed to a "memory integrated circuit" [Nishiguchi, 1:7]. It contains an on-chip programmable counter for delaying the read out operation until the number of clock cycles following a read request equals the programmed count. The invention is described using a ROM type of memory as an example. As expressly stated in the specification, it is obvious that the invention disclosed is "equally applicable in the case of a RAM" [Nishiguchi, 1:12]. The memory device includes "a memory section 6 including memory cells in a matrix of rows and columns" [Nishiguchi, 1:17].

5.15.1.1.2. The memory integrated circuit of Nishiguchi includes circuitry to receive an external clock signal, which is "commonly utilized both for the CPU and this ROM" [Nishiguchi, 3:36]. Such clock signals are of a fixed frequency. The circuitry to receive the external clock is not specified in Nishiguchi, but must be present and can simply be an input buffer.

5.15.1.1.3. The memory integrated circuit of Nishiguchi includes a programmable access-time register in the form of the selection circuit 19 comprising cross-points 191 to 196 shown in Fig. 2. The state (open or closed) of each of the cross-points in selection circuit 19 may be programmed in the same way the ROM memory itself is programmed [Nishiguchi: 4:40-43; 6:17-19; 6:64-68]. At least until so programmed, they are programmable. An access delay time is specified in numbers of clock cycles by programming one of the cross-points to the closed position. For instance, if cross-point 193 is programmed to a close position, the access delay time will be 3 clock cycles. Access delay time can be changed by reprogramming the cross-points [Nishiguchi, 4:31-34]. The memory device responds to a CPU read request by setting a ready flag after the programmed number of access delay clock cycles have transpired [Nishiguchi, 3:35; 4:35-40; 5:34-38].

5.15.1.2. German Patent Publication DE 37 42 487, Published July 1988, and English Translation ("Kawai")

5.15.1.2.1. Kawai discloses a variable delay circuit buffer, which is a type of FIFO buffer in which a timing relationship is specified between the data input into the FIFO and its output. FIFO buffers are semiconductor memory devices. The buffer contains a core memory cell array made of rows and columns [Kawai, Fig. 3]. The memory cells can be static memory cells (SRAM core) or dynamic memory cells (DRAM core).

5.15.1.2.2. Kawai discloses a synchronous buffer device that is clocked by an external clock signal. The external clock signal is labelled as ϕ in Fig. 3. Receiver circuitry such as a buffer is used to receive the external clock signal. Kawai discloses a register corresponding to an amount of delay that is programmed using a select signal bus of k bits [Kawai, Fig. 2]. The buffer device reads in data, which it holds for a time equals to the variable delay value programmed into the access-time register of the buffer, and then outputs it onto a bus. On rising of the L-th clock ϕ , ... the initial data written in is output to the DO1-DOm ... The L is a value that is programmable and externally set by users.

5.15.1.2.3. Kawai shows that the data output circuit 6 includes a plurality of output drivers for outputting data on an m-bit wide bus (DO1-DOm). Each m-bit wide data is outputted in response to an implicit read request (set at the time the data is clocked into the buffer), and only after the programmed number of clock cycles has transpired (as discussed above). The data output is done synchronously with the external clock signal. Kawai discloses a control circuit 7, coupled to the circuitry that receives the external clock signal ϕ , that generates a control signal for timing the data output operation [Kawai, Fig. 3].

5.15.1.3. Japanese Patent Application JP-A-01-236494, Published September 21, 1989, and English Translation ("Akimoto")

5.15.1.3.1. Akimoto claims "a semiconductor integrated circuit device characterised by comprising a delay circuit whose delay time changes in accordance with a selection control signal that is received via a prescribed pad or external terminal" [Akimoto, claim 1]. Akimoto discloses a clocked SRAM with an on-chip programmable timing generator circuit TG [Akimoto, Fig. 4]. That circuit receives inter alia an external clock CK and delay setting signals SS0-SS1, and generates three internal clocks ϕ_{ce} , ϕ_{we} and ϕ_{oe} for clocking the address, input, and output latches respectively. Each of the address capture, input or output clocking operation can be delayed from the external clock CK by use of the delay setting signals SS0-SS1, as shown in exemplary fig. 1 for the write clock ϕ_{we} .

5.15.2. Automatic Precharging (Claim 2)

5.15.2.1. US Patent 4,754,433, Issued June 28, 1988 ("Chin")

5.15.2.1.1. Chin describes an asynchronous DRAM having multiplexed twin I/O line pairs [Chin, 3:20-27]. Chin utilises a sense amplifier multiplexing scheme in which the DRAM alternates between sets of sense amplifiers in a dual bit per CAS mode of operation [Chin, 3:60-4:41]. Whilst one set of sense amplifiers is read, the other set is precharged in preparation for the next read [Chin 7:66-8:7].

5.15.2.2. US Patent 4,845,677, issued July 4, 1989 ("Chappell")

5.15.2.2.1. Chappell discloses a self-timed DRAM [Chappell, 4:29-34; Fig. 3-2]. The DRAM is internally partitioned into blocks. The internal partitioning of the chip allows one block of the memory to be precharged immediately after a read operation, in preparation for the next read operation, while data is read from another sub-array [Chappell 5:28-31]. A wait state generator inserts wait cycles if one of the three subsequent read operations requires access to the same block and thus must wait for precharge of that block to be completed [Chappell, 5:46-61 and Fig. 7b]. As shown in Fig. 7b, where a read operation to one block (block D) is immediately followed by a read operation to the same block, three wait cycles are inserted. The number of wait cycles inserted is enough to ensure that the precharge of the block to be read is complete [Chappell, 5:58-61].

5.15.3. It was conventional on the priority date for DRAMs to precharge the sense amplifiers and hence the bit lines automatically after a normal mode memory access operation in preparation for a subsequent read request from another row. This is acknowledged in the Patent [Patent, 7:30-34]. It follows that the subject-matter of claim 2 is obvious.

5.15.4. Block Size Information (Claim 4)

5.15.4.1. US Patent 4,763,249, Published August 9, 1988 ("Bomba")

5.15.4.1.1. Bomba discloses a system that includes a plurality of bus devices interconnected by a synchronous, multiplexed bus. The bus device can be constructed as a memory device with a plurality of storage locations and interconnection circuitry [Bomba, abstract]. The interconnection circuitry forms an integral part of the memory device [Bomba, 6:63-65]. A master clock 144 connected to the communications path generates clocking signals (time and phase signals) for the bus devices [Bomba, 9:14-32 and Figs. 3A and 3B]. These signals are received by, inter alia, the interconnection circuitry of the memory device. A two-bit data length code is placed on the highest bits of data lines D[31:30] during a command/address cycle of a read transaction [Bomba, 13:59-14:4 and Fig. 4A]. The lower 30 bits contain the device "address" (the address is the 30 bit storage location where the transaction is to take place) [Bomba, 13:59-14:4 and 14:13-15]. The data length code specifies the length of the data transfer that is to take place, e.g. one to four cycles of 32 bit data [Bomba, 15:18-28]. The operation code for the read command is transmitted over information lines I[3:0] at the same time [Bomba, 13:59-61]. As can be seen from Figs. 1A, 2 and 3B of Bomba, each of these signals is received by, inter alia, the interconnection circuitry of the memory. A memory device on the bus, containing the relevant address, confirms receipt of the command/address cycle and may begin the data transaction [Bomba, 15:49-63]. The memory device outputs data on the data lines D[31:0] during a first data cycle. The memory device continues to output new data for as many data cycles as are specified during the command/address cycle [Bomba, 15:29-16:58]. Data is generally placed on the data lines at the leading edge of internal clock TCLK [Bomba, 9:51-54]. TCLK is the transmitting clock generated locally in the memory device from the time and phase components of master clock 144 [Bomba, 8:46-63 and 9:27-39 and Fig. 3B].

5.15.4.2. US Patent 4,394,753, Published July 19, 1983 ("Penzel")

5.15.4.2.1. Penzel describes a highly integrated memory module. As illustrated in Fig. 1 of Penzel, the highly integrated memory module is a semiconductor memory device that includes a memory cell array (DECODER MEMORY) arranged as a plurality of rows and columns [Penzel, 2:45-53]. As illustrated in Fig. 5 of Penzel, during a chained memory access (block read/write), the column address strobe signal CAS_i is pulsed [Penzel, 6:44-50]. Fig. 5 shows CAS_i as being a fixed frequency signal. The memory device of Penzel includes a mode register [Penzel, 3:9-11]. The mode register is a nine bit register in which two bits M0, M1 determine whether the read is a '1 read, a '4 read or a '8 (or '9) read [Penzel, 3:37-42]. Two further bits, M2, M3 determine the number of chained accesses in a read operation (block size) [Penzel, 3:42-44 and 3:46-57]. The mode register is programmed externally via package pins P0-P8 [Penzel, 3:63-65]. The block read operation is described [Penzel, 4:21-26; 6:34-65]. As is clearly illustrated in Fig. 5, successive data is output in response to successive pulses of the external clock signal CAS_i.

5.15.4.3. IEEE Standard for a Simple 32-Bit Backplane Bus: NuBus - ANSI/IEEE Std 1196-1987 ("NuBus")

5.15.4.3.1. The NuBus standard is a synchronous computer backplane bus standard in which the bus is used to connect devices and to provide certain

resources to the connected devices [NuBus, Fig. 1]. Conventional memory modules are one type of device that may be attached to the NuBus [NuBus, p. 42 (A.1 Note)]. The NuBus modules receive an external clock source (central system clock) [NuBus, Fig. 1; p. 4 (Section 2.1.1)]. The clock source is common to each NuBus module. The clock is a fixed frequency clock, nominally operating at 10 MHz, and is used to synchronise bus arbitration and data transfers [NuBus, p. 4 (Section 2.1.1)]. The clock signal is driven from one end of the bus to termination at the other end [NuBus, p. 44 (Section A.7)]. NuBus provides for block data transfers in block sizes of 2, 4, 8 and 16 words [NuBus, p. 11 (Section 3.1.4)]. The number of data words transferred is controlled by the master and communicated during the start cycle. The block size and block starting address is transmitted over the 32 bit multiplexed address and data lines (which run between the master and slave devices) while the START* signal is asserted [NuBus, Fig. 4; pp. 11-13]. The slave device drives the first word of the requested data onto the 32 bit multiplexed address and data lines. The start of the data transfer is synchronous with respect to CLK* (occurring at the rising edge) [NuBus, Fig. 5; p. 12]. Blocks of data are output from the slave synchronously with respect to CLK* until the desired block size is reached [NuBus, Fig. 5].

5.15.4.4. Scalable Coherent Interface ("SCI")

- 5.15.4.4.1. "The Scalable Coherent Interface Project (SuperBus)", SCI-22Aug88-doc ("SCI A")
- 5.15.4.4.2. "Scalable Coherent Interface", SCI-28Nov88-doc20 ("SCI B")
- 5.15.4.4.3. P1596: "SCI, A Scalable Coherent Interface", SCI-28Nov88-doc 2 ("SCI C")
- 5.15.4.4.4. "Proposal for Clock Distribution in SCI" - 5/5/89 ("SCI D")
- 5.15.4.4.5. Norsk Data Report - "A Proposal for SCI Operation" by Knut Alnes - November 1988 ("SCI E")
- 5.15.4.4.6. "Scalable I/O Architecture for Buses" by David V. James, SCI-28Nov88-doc3 ("SCI F")
- 5.15.4.4.7. SCI is an interface standard used to connect devices, including processors, I/O devices and memory [SCI A, p. 6]. A device can be a simple memory module [SCI F, Fig. "Board Architecture" p. 3]. SCI uses a 16 bit wide synchronous, packetized bus, to carry address, data and control information [SCI A, p. 4]. Each SCI device receives a signal from a central system clock [SCI A, Fig. 3; SCI B, p. 2]. An SCI device receives a request packet with target, source, control and address information [SCI A, Figs. 4 and 15]. SCI supports operations of 32, 64, 128 and 256 data blocks, and 1-16 byte subsets of the 16-byte block [SCI A, p. 8; Fig. 11; SCI C, p. 11]. The block size information is conveyed in the transfer code, which forms part of the control information received by the device [SCI A, Figs. 14 and 15; p. 14; SCI C, p. 13 (Header Command)]. The targeted device outputs a variable size data block (i.e., Data Word 0 to Data Word n) in accordance with the block size information [SCI A, Fig. 15; p. 14]. All inputs and outputs on the bus are synchronous with respect to the external system clock [SCI A, p. 3]. All devices in SCI operate synchronously with respect to the system clock.

5.15.4.5. US Patent 4,785,394, Published November 15, 1988 ("Fischer")

5.15.4.5.1. Fischer describes a multiprocessor computer system arranged around a split-transaction bus [Fischer, abstract]. Bus devices are notionally classified into "initiators," including CPU modules, and "responders," including memory modules [Fischer, Fig. 1]. The split-transaction bus carries all information between bus devices and no point-to-point signals are used [Fischer, 7:36-40]. Each of the memory modules includes conventional memory components [Fischer, 6:6-7]. Some or all of the memories in the modules may be cache memories [Fischer, 6:17-19]. Both the conventional memories and cache memories are semiconductor memories. Each includes semiconductor logic components [Fischer, 5:47-49]. A typical responder module is illustrated in Fig. 2 (right hand side) of Fischer. As can be seen, it includes clock receivers 76 that receive clock signals B.CLK 0, B.CLK 1. As shown in Figs. 3A and 3B of Fischer, these clock signals are fixed frequency square waves in quadrature [Fischer, 8:53-55]. They are generated by a backplane clock generation circuit 74. Clock skew is eliminated by the use of equal length clock signal conductors, ensuring reliable synchronisation [Fischer, 8:33-50]. The initiating transaction of a read operation is illustrated in Fig. 6B [Fischer, 10:32-35]. The format of the address information transmitted in bus cycle X of Fig. 6B is shown in Fig. 7A [Fischer, 14:26-29]. It includes three fields. A first field 86 is a two-bit field indicating the nature of the transaction (read, test and set, scrub or write) [Fischer, 13:59-64]. A second field 88 is a 28 bit memory address that identifies the responder to which the request is addressed and the address within that responder [Fischer, 13:64-14:4]. A third field 90 indicates whether one, two or four doublewords are to be transferred in response to the request [Fischer, 14:4-8]. Each clock receiver 76, illustrated in Fig. 4 of Fischer, generates four internal clock signals B0 ... B3, one quarter cycle out of phase with each other [Fischer, 9:6-25; Figs. 3C-3F]. The format of data placed on the bus in response to a read request is illustrated in Fig. 7B. As shown in Fig. 6B, one doubleword is output per bus cycle. The Fischer system is synchronous [Fischer, 8:33-50; 9:26-38]. Control, address and data signals transferred from bus master to bus slave (initiator to responder or vice versa) are designated as signals B.DAT31-0 [Fischer, 15:43-50]. These signals are asserted on the rising edge of B0 and negated on the rising edge of B3 (three quarter-cycles later). B0 is synchronised with external clock signal B.CLK 1 [Fischer, Figs. 3A, 3C and 4].

5.15.4.6. US Patent 4,315,308, Published February 9, 1982 ("Jackson")

- 5.15.4.6.1.** Jackson discloses a system with a microprocessor attached to a Bus Interface Unit (BIU) which provides for interface control of data transfers between the processor and devices, including memory devices [Jackson, 4:17-21]. Each unit in the system receives a clocking signal CLKA to control synchronisation [Jackson 4:36-44; 5:36-38]. CLKA is of fixed frequency [Jackson, Fig. 2].
- 5.15.4.6.2.** Jackson discloses block modes of up to 20 bytes per read request [Jackson, 3:16-19; 6:4-6]. Specifically, Jackson supports transactions of 1, 2, 4, 6, 8, 10, 16 and 20 bytes [Jackson, Fig. 3]. A three bit long sequence in a control specification is used to specify block size [Jackson, Fig. 3 (bits 10, 11 and 12); 5:29-30]. The memory returns the requested number of data bytes to the BIU [Jackson, 6:8-12]. As the memory continues to output the data, the BIU buffers and aligns the data and then transfers the data across the ACD bus to the processor. This process is repeated until the processor has received the requested number of bytes [Jackson, 6:20-23]. A read operation takes place over a

plurality of clock cycles [Jackson, 5:38-39; Fig. 9]. The block size information is a binary representation of the amount of data to be output [Jackson, Fig. 3].

5.15.4.6.3. Jackson discloses the use of a control specification whereby the operation code (read request), block size information and the eight least significant bits of the address are placed on the bus at the same time [Jackson, Figs. 2 and 3; 5:53-58].

5.15.4.7. Japanese Patent Application No. S63-142445, Published June 14, 1988, and English Translation ("Taguchi")

5.15.4.7.1. Taguchi discloses a memory device having an array formed from a plurality of memory cells [Taguchi, Fig. 1 in conjunction with Figs. 5 and 6 and 4:27-29]. A data length register holds the entire size of the data to be accessed [Taguchi, p. 4]. In addition, Taguchi discloses the use of a block length register to define the size of each block and a data length register to define the number of blocks that are to be accessed in an operation [Taguchi, p. 4]. Block data of the size specified in the block length register is read from the memory, the next block of memory to be accessed is calculated from the starting address plus the block size, another block read is performed, and the new starting address is calculated [Taguchi, p. 4]. Counters monitor the number of block reads left in the block access operation. The reading process continues until the amount of data specified in the data length register is output [Taguchi, p. 4].

5.15.4.8. Microprocessor Report, "ECL Bus Controller Hits 266Mbytes/s; MIPS R6020 Handles CPU, RAM, I/O Interface", Published January 24, 1990, ("MIPS")

5.15.4.8.1. MIPS R6020 describes a bus architecture designed to maximise performance of the MIPS R6000 CPU [MIPS, p. 12, col. 1:1-5]. Each bus request is packetized and includes control information in a command word [MIPS, p. 13, col. 1:41 - col. 2:5]. The command word includes address information and information identifying whether the request is a read or a write and whether it is a single word or a block transfer [MIPS, p. 13, col. 2:2-5]. The request is completed by a second word following the command word, including the balance of the address information [MIPS, p. 13, col. 2:6-8].

5.15.5. Delay Locked Loop (Claim 5)

5.15.5.1. At the priority date of the Patent, DLLs were common general knowledge. DLLs enable better alignment of regularly cyclic digital signals at the same frequency.

5.15.5.2. UK Patent Application GB-2,197,553, Published May 18, 1988 ("Lofgren")

5.15.5.2.1. Lofgren describes a digital phase locked loop circuit (DLL) [Lofgren, abstract]. One principal application of the DLL described is to provide "optimum timing for control of high speed dynamic RAM devices" [Lofgren, 1:14-18]. Lofgren discloses the use of two identical delay lines. One delay line 12 is clocked by a local oscillator 20 and used to calibrate the DLL. The other delay line 18 is used to provide an accurate delay to an input signal [Lofgren, 1:130-2:5].

5.15.5.2.2. Each delay line 12, 18 consists of plural delay elements D_n , the amount of delay introduced by each of which being controlled by the level of two control signals VCP, VCN, generated by a charge pump 16 in response to phase errors detected in the first delay line 12 [Lofgren,

3:78-114; 4:23-33; 4:62-68; 5:78-96]. The delay introduced by delay line 12 is varied until it is exactly one period of the oscillator 20 [Lofgren, 4:104-106]. Identical control signals are applied to the second delay line 18. It is then known that the delay elements D_n in the second delay line 18 will introduce one n^{th} of the period of the local oscillator. The second delay line is a multi-tap line, allowing one of a plurality of delays to be selected [Lofgren, 2:5-13].

5.15.5.3. IEEE Journal of Solid State Circuits, Vol. 25, No. 1, February 1990, "An On-Chip Smart Memory for a Data-Flow CPU" ("Uvieghara")

5.15.5.3.1. Uvieghara describes a high performance substrate CPU that has an embedded smart memory of the type known as a "register alias table" ("RAT"). The RAT is a multi-port content-addressable memory supporting branch prediction and exception handling. An experimental 1240 bit RAT is described [Uvieghara, abstract]. The RAT is a synchronous DRAM [Uvieghara, Fig. 5] having a block read mode [Uvieghara, 87:2-4].

5.15.5.3.2. The RAT uses a PLL-based clock generator on-chip [Uvieghara, Fig. 13; and p. 92 "All clocks are generated by a ... PLL-based clock generator."]

5.15.5.4. Scalable Coherent Interface ("SCI")

5.15.5.4.1. SCI discloses the use of an internal clock signal generated by a digital phase lock loop to eliminate clock skew [SCI B, p. 2; SCI D, pp. 1-5]. The internal clock is used for outputting data on the external bus [SCI B, p. 2]. The digital PLL disclosed is a DLL: it shows the use of a delay line and no VCO.

5.15.5.5. US Patent No. 4,637,018, Issued January 13, 1987, ("Flora")

5.15.5.5.1. An example of a DLL is described in Flora. The DLL of Flora is designed to align the outputs of a distributed clock driver circuit with an accurately delayed external clock signal [Flora, 3:6-16]. As the device is a clock distribution system, exact alignment with the external clock signal is not essential, but alignment of the outputs of several chips with each other is [Flora, 3:22-28]. For this reason the external clock signal receiving line is bifurcated. One branch is subject to an accurate delay and the other includes a multi-tap delay line introducing a variable delay to the clock driver circuit. The outputs of the clock driver circuitry are aligned with the delayed external clock [Flora, 3:60-68]. However, if accurate phase alignment with the external clock were required, the delay line would be removed, or re-introduced in the same branch as the variable delay to give an accurate delay of just short of one clock cycle. In this way, the outputs of the clock driver circuitry would be accurately aligned with the external clock. This all results from the application of common general knowledge at the priority date.

5.15.5.6. Japanese Patent Application JP-A-01-284132, published November 15, 1989, and English Translation ("Kosugi").

5.15.5.6.1. Kosugi discloses a digital phase locked loop (DPLL) that is used for aligning the internal read clock 3 (clocking the output) of a memory device 1 with its internal write clock 2 (clocking the input). The memory write clock 2 is itself generated from an external clock.

5.15.5.7. Motorola MC88200 Cache/Memory Management Unit User's Manual, Published 1988 ("MC88200")

5.15.5.7.1. The MC88200 is a single chip synchronous device that contains high-speed cache memory. The MC88200 has a large number of on-chip, dynamically programmable configuration registers, as shown in table 1-1. The device generates on-chip all internal timing signals from an external clock signal CLK. The MC88200 internal clock is normally phase locked to the external clock signal CLK in order to minimise the skew between the external and internal signals [MC88200, page 4-9].

5.15.5.8. Johnson et al., "A Variable Delay Line PLL for CPU-Coprocessor Synchronization," IEEE Journal of Solid-State Circuits 1988 ("Johnson I")

5.15.5.9. European Patent Application EP 329,418, Published August 23, 1989 ("Johnson II")

5.15.5.10. Johnson et al., "A Variable Delay Line PLL for CPU-Coprocessor Synchronization," ISSCC 88, February 18, 1988 ("Johnson III")

5.15.5.10.1. Johnson I, Johnson II and Johnson III disclose an integrated phase-locked loop (PLL) for preventing bus contention between a CPU and its floating-point co-processor (FPC), both of which are connected to the same cache data bus [Johnson I, Fig. 1; Johnson II, Fig. 1; Johnson III, Fig. 1]. The CPU and FPC are clocked by the same system clock CLKIN [Johnson I, Fig. 2; Johnson II, Fig. 2; Johnson III, Fig. 2]. The CPU and FPC output timing signals OE are internally generated from that clock. Ideally, the two OE signals are complementary of each other, so that only one chip seizes the bus at any one time, the other chip being disconnected from the bus by assertion of its OE signal. However, differences in clock wiring length, both on-chip and off-chip, can create substantial misalignment or skew in the chips' output timing signals [Johnson I, 1219].

5.15.5.10.2. A PLL is implemented in the FPC to re-align the OE output timing signal of the FPC chip with the OE output timing signal of the CPU [Johnson I, Fig. 2; Johnson II, Fig. 2; Johnson III, Fig. 2]. The phases of the CPUOUT and FPCOUT signals, corresponding to the OE signals of the CPU and FPC respectively, are first compared with each other to generate an error signal [Johnson II, 3:59-4:28]. That error signal is then used in a feedback loop to adjust the delay, i.e., phase, of the FPC output timing signal [Johnson II, 4:29-51].

5.15.5.10.3. The PLL does not use a traditional voltage-controlled oscillator (VCO). Instead, it uses a voltage-controlled delay line (VCDL) [Johnson I, 1219-20, Johnson II, 37-58]. Hence, the "PLL" disclosed in Johnson I, Johnson II and Johnson III is in fact a DLL.

5.15.5.11. It was obviously desirable on the priority date of the Patent that internal device clock signals should be aligned as closely as possible with the external clock signal. For example, the further internal clock signal ϕ_c of Yamaguchi should be aligned with the external clock signal SC [Yamaguchi, 26:17-19]. The objective problem to be solved by a device according to granted claim 5 vis-à-vis other devices is the provision of more accurately aligned internal clock signals. This problem is solved by using a DLL, as would have been well known to a person skilled in the art at the priority date.

5.15.5.12. It follows that the subject-matter of claim 5 is obvious.

5.15.6. Synchronous Input Sampling (Claim 10)

5.15.6.1. Synchronous device interfaces were well known on the priority date as a mechanism for improving high-speed operation of semiconductor memory devices.

5.15.6.2. US Patent 4,631,659, Published December 23, 1986 ("Hayn II")

5.15.6.2.1. Hayn II discusses the advantages of a synchronous memory interface in the following terms: "Using a synchronous interface to an external memory saves time and improves the performance of the microcomputer system." [Hayn II, 1:32-34].

5.15.7. Multiplexed Bus Architecture Using Request Packets and Device IDs (Claims 9 and 11-16)

5.15.7.1. Multiplexed buses using request packets were common general knowledge at the priority date of the Patent.

5.15.7.2. US Patent No. 4,763,249, Issued August 9, 1988 ("Bomba")

5.15.7.2.1. Bomba discloses a bus interface device 18 for use with a synchronous, multiplexed bus. The bus transmits device ID, burst length, address and data in packetized form over the same bus lines [Bomba: 8:19:26]. A bus interface device is "located integrally within" each memory device 12 [Bomba, 6:64]. Each memory device receives a plurality of external clock signals and generates local transmit and receive clocks. Each memory device contains a large number of configuration registers (200-216) that are used for programming the device ID, bus arbitration scheme, etc onto the memory device, typically on system power-up or during an initialisation sequence [Bomba, 22:6:44].

5.15.7.3. Scalable Coherent Interface ("SCI")

5.15.7.3.1. Each SCI node has a unique 16 bit identification code stored in a register [SCI A, p. 6; SCI E, pp. 1-4]. The ID information is sent in the request packet and received by the SCI nodes [SCI A, pp. 6-7; Fig. 9]. SCI nodes determine whether the target identification code in the request packet matches the identification code of the node [SCI E, pp. 1 and 5-6]. The targeted node outputs a variable size data block (i.e., Data Word 0 through Data Word n) in accordance with the block size information [SCI A, Fig. 15; p. 14]. All inputs and outputs on the bus are synchronous with respect to the external system clock [SCI A, p. 3]. All nodes in the SCI operate synchronously with respect to the system clock. The identification register contains a Global ID and Local ID. The Global ID is the 10 MSB of the node ID and is used to identify a ring on the SCI network (which contains several nodes) [SCI E, p. 1]. Node ID's are assigned during an initialisation sequence after reset of the SCI. A hardware or software based protocol is employed to assign node identifications. Following the assignment of node identification, the SCI master will initialise the interface registers of the various nodes and will cause the node ID to be written into the individual registers [SCI E, pp. 2-4]. Thus, SCI discloses that "the programmable identification register is programmed after power is applied to the memory device duration initialisation of the memory device" and that "the programmable identification register stores an identification value during an initialisation sequence of the memory device."

5.15.7.4. US Patent 4,785,394, Published November 15, 1988 ("Fischer")

5.15.7.4.1. As discussed above, the format of the address information transmitted in bus cycle X of Fig. 6B is shown in Fig. 7A [Fischer, 14:26-29]. It includes three fields. The second field 88 is a 28 bit memory address that identifies the responder to which the request is addressed and the address within that responder [Fischer, 13:64-14:4]. No point to point signals are used [Fischer, 7:36-40]. Thus, each memory module must decode the MSBs of the address to determine whether to respond. Accordingly, the module must contain an internal address space identification with which to compare the MSBs. That address space identification must be programmable, and in Fischer it is received from the backplane [Fischer, 16:21-27]. This must be done during an initialisation sequence following power-up or the memory would be unable to respond to requests.

5.15.7.5. US Patent No. 4,481,572, Issued November 6, 1984 ("Ochsner")

5.15.7.5.1. Ochsner discloses a time-shared bus for use with digital computers. The computers are structured around a high-speed bi-directional time-shared bus in which both address and data are carried on the same lines. Each unit on the bus has associated with it a response time represented by a predetermined number of time slots between its receipt of a request for information and its providing of information onto the bus. That number of time slots' delay is resident in both a respective unit controller and in the resource controller [Ochsner, abstract].

5.15.7.5.2. The bus is driven by an 82 nanosecond clock [Ochsner, 5:56-62]. The resource controller is clocked at twice the rate of the bus [Ochsner, 5:62-68]. The memory delay times are expressed in terms of the number of half time slots or bus cycles so that the memories may be operated at the highest speeds practical. Because the clock speed is known, the delay of each memory in terms of the number of 42 nanosecond increments is a predetermined characteristic of memories M1 and M2 and generally will be hard wired into the resource controller 22 [Ochsner, 6:5-20]. The system is synchronous [Ochsner, 5:1-5]. Individual memories are identified by signals on lines 28 and 30 [Ochsner, 6:9].

5.15.7.6. Microprocessor Report, "ECL Bus Controller Hits 266Mbytes/s; MIPS R6020 Handles CPU, RAM, I/O Interface", Published January 24, 1990, ("MIPS")

5.15.7.6.1. MIPS R6020 describes a bus architecture designed to maximise performance of the MIPS R6000 CPU [MIPS, p. 12, col. 1:1-5]. The bus is a relatively narrow bus, and this is achieved by multiplexing address, data and control information onto the same bus lines [MIPS, p. 12, col. 2:11-19].

5.15.7.6.2. Dedicated chip select lines are not required, owing to a daisy-chained reset signal that allows each device on the bus to determine its slot number [MIPS, p. 12, col. 2:29 - p. 13, col. 1:17]. All requests are latched by all devices on the bus [MIPS, p. 13, col. 1:31-40].

5.15.8. Low Voltage Swing (Claim 17)

5.15.8.1. A "low voltage swing" signal was a term of art at the priority date. It was used in the Grandparent Application without further explanation.

5.15.8.2. IEEE 1987 CUSTOM INTEGRATED CIRCUITS CONFERENCE, "Self Termination Low Voltage Swing CMOS Output Driver", Tom Knight and Alex Krynn, CH2430-7/87/0000-0289-0000-0292 ("Knight")

5.15.8.2.1. An example of a low voltage swing output driver is described in Knight. The advantages of using low voltage swing signals were common general knowledge, including low power consumption and higher inter-chip communications speeds [Knight, 289, "Introduction"].

5.15.8.3. Ogiue et al. "13-ns, 500 mW, 64-kbit ECL RAM Using HI-BICMOS Technology," IEEE Journal of Solid-State Circuits 1986 ("Ogiue")

5.15.8.3.1. Ogiue discloses a bipolar CMOS SRAM memory device that uses ECL low voltage swing signalling for its I/O signal interfaces to achieve faster access time. Each ECL low voltage swing signal input is internally amplified to full MOS signal level by using a level shifter buffer circuit of the type shown in Fig. 6.

5.15.8.4. The objective problem solved vis-à-vis known devices by a device according to granted claim 18 is the reduction in power consumption or the increase in inter-chip communications speeds. It would be entirely natural to seek to realise those advantages in any of the memory systems previously discussed. The problem is solved by using low voltage swing signals, as would have been well known to a person skilled in the art at the priority date.

5.15.8.5. It follows that the subject-matter of claim 17 is obvious.

6. INSUFFICIENCY ART. 83 EPC

6.1. The invention claimed in granted claim 1 is insufficiently described.

6.1.1. Only One Selectable Access Time

6.1.1.1. As discussed above, granted claim 1 appears to cover devices in which, if a specific value is stored in the programmable access time register, the device delays a predetermined number of clock cycles and outputs data synchronously with respect to the external clock signal; and in which, if another value is stored, the device does not delay.

6.1.1.2. There is no disclosure in the Patent of how to construct a semiconductor memory device that operates in this way. All that is described is a device that always delays for a number of clock cycles dependent upon the value stored in the access time register [Patent, 7:17-18]. It is not apparent from the Patent how this gap is to be bridged. Based upon the information contained within the Patent, the person skilled in the art would be unable to practise the invention throughout the range claimed in granted claim 1.

6.1.1.3. It follows that the subject-matter of claim 1 is insufficiently described contrary to Art. 83 EPC.